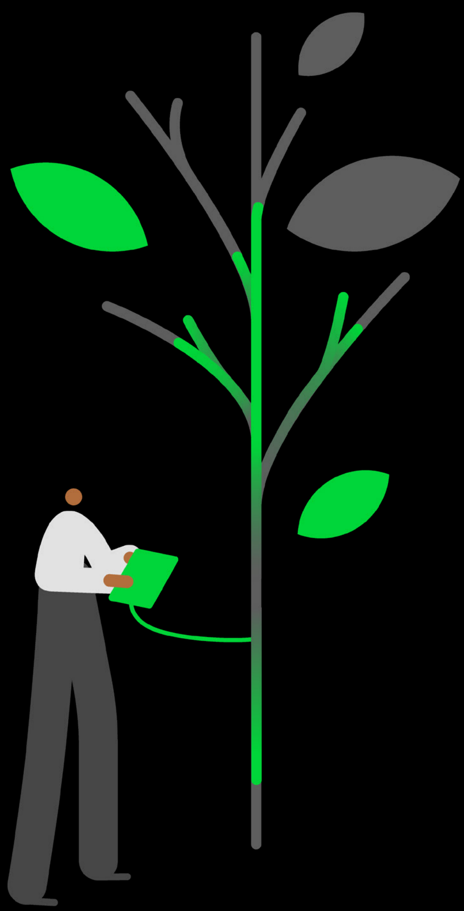

The role of LCAs in decarbonising the cloud

As organisations scale their digital transformation efforts, one area often overlooked is the environmental impact of software and the associated cloud infrastructure that supports this.

Sage

The use of software has increasingly been migrating to the cloud as it is, traditionally, seen as a more energy-efficient alternative to on-premise datacentres. However, the cloud has increasingly become a significant source of global greenhouse gas (GHG) emissions and we now face a dilemma between technology and sustainability. Cloud-based services have brought significant benefits, such as cost savings, accessibility, and data security(1), but its growing environmental impact demands our attention.

In this article, we explore the challenge of quantifying GHG emissions from software, and its associated infrastructure, by looking at the role of different carbon lifecycle assessments (LCAs) in supporting and promoting transparency, and how identifying emission hotspots through these assessments can drive impactful carbon reduction actions for software manufacturers, cloud providers, and the consumers of these services.

The Rising Emissions of Cloud Computing

In the early days of cloud adoption, the primary appeal was its ability to scale operations without the need for costly, energy-intensive, on-premise infrastructure. Cloud migration promised cost efficiency, flexibility, accelerated growth, and the possibility of reducing an organisation's carbon footprint by leveraging larger, more efficient, shared datacentres. It also reduced concerns over end-of-life disposal relating to on-premise infrastructure. Thanks to economies of scale, hyperscale datacentre providers like Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform can operate more sustainably than businesses with their own infrastructure.

The shift to cloud computing has improved the efficiency of Information and Communication Technology (ICT) in a way that parallels Moore's Law(2) – the principle that computing power doubles roughly every two years while costs decrease.

Similar to Moore's observation of exponential growth in processing power vs a reduction in the physical footprint and energy costs of devices, cloud computing has made ICT more efficient by centralising resources, optimising hardware usage, and improving energy efficiency at scale. For decades the energy required to do the same amount of computation has halved every two-and-a-half years, a trend known as Koomey's law(3). While this remains true, the exponential growth of cloud computing has led to an explosion in energy demand and the overall growth in demand for ICT has surpassed efficiency improvements, resulting in a rise in total impacts. Today, the carbon footprint of ICT is estimated at between 1.4-4% of global GHG emissions, and rising fast (4; 5; 6; 7).

The International Data Corporation (IDC) projects that global data centre electricity consumption will more than double by 2028, reaching 857 TWh(8). However, the environmental impact of data centres goes beyond energy use; the lifecycle of IT hardware, network infrastructure, and associated cooling systems also play a critical role.

As an industry we are now grappling with the paradox of cloud adoption: it fuels innovation and drives economic growth while, in some instances, also exacerbating emissions that contribute to the climate crisis. It brings efficiency to the global economy, but this, too, can backfire into rising global emissions.
This tension is compounded by the rapid adoption of artificial intelligence (AI), machine learning models, blockchain, and other high-performance computing tasks – activities that place increasing demand on energy, water and natural resources.

This raises a key question: How does the industry make cloud operations more efficient and ensure that the footprint of cloud services is aligned with global climate goals?
One starting point to this debate is recognising the need to transparently, realistically, and consistently, measure GHG emissions related to the cloud, and the software that drives this demand. How, and under what circumstances, can the measurement of GHG emissions from ICT be meaningful and understandable to the industry?

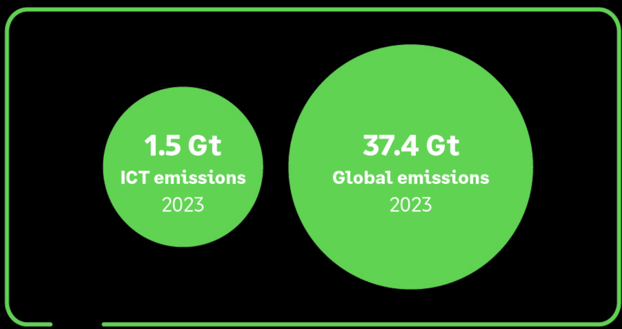

The Role of Carbon Lifecycle Assessments in ICT Sustainability

Carbon lifecycle assessments (LCAs) are tools that provide a comprehensive view of the environmental impact of a product, service, or system throughout its entire lifecycle – from cradle to grave. In the context of software development, LCAs enable us to calculate the cumulative emissions across each life-cycle stage of the software service delivery process, such as the software development by developers, writing of the code, the datacentres, the networking infrastructure, the servers, the energy mix that powers them. A key aspect of these emissions includes the (length of) time, and frequency of software use, as well as how the software is used.
This helps map and identify the significant emissions sources, or 'hotspots', within the system boundary which, in turn, facilitates decarbonisation actions.

One of the challenges to improve the sustainability of ICT is the lack of standardised metrics. This makes it difficult for users to understand the true environmental cost of the services they consume, including the impact of historic software development. Most of the attention on the GHG impact of ICT has tended to be focussed on individual life-cycle stages, in particular on the energy efficiency (such as the Power Usage Effectiveness (PUE)) and Scope 1, 2 and 3 emissions of data centres. Applying LCA methodology to software products can help achieve a more holistic understanding across the entire life-cycle.

As LCAs can be used to disclose the full carbon impact of ICT services, they can be used to guide sustainability strategies, moving beyond energy efficiency to embrace circular economy principles like hardware recycling, extending product lifespans, and promoting sustainable software design practices. Activities that will facilitate the decarbonisation of the industry.

There are two basic methodological tools for life cycle analysis. They have complementary strengths and weaknesses, so it is possible to combine them into a hybrid 'method' that mitigates the weaknesses that are inherent in each method when used on its own. Within any given resource, careful blending of the two approaches generally enables the most realistic overall result. However, any output centres on decisions about attribution and background assumptions.

The first, and most widely understood of these methods, is process-based life cycle analysis (P-LCA). This is considered a 'bottom-up' approach because it aims to understand the carbon impact of each stage in the life cycle of a product by mapping out the production process and obtaining physical emissions data. This means that data on the energy consumption and physical items used within each stage must be gathered to calculate embodied and use-phase emissions.

> **A 'bottom-up' approach aims to understand the carbon impact of each stage in the life cycle of a product by mapping out the production process and obtaining physical emissions data.**

This approach can be extremely resource-intensive because it requires gathering and analysing large amounts of data for each lifecycle stage. Comprehensive data often does not exist for each stage, meaning that a P-LCA can be prone to uncertainties that risk creating an incomplete map of emissions. This level of granularity may not align with the priorities or resources of organisations looking to implement carbon reduction actions, especially in dynamic, fast-changing sectors like software development and the ICT industry in general. However, at its best P-LCA can provide specific and detailed assessments of the critical hotspots in a life cycle.

One major limitation with P-LCA is the judgement required in determining the study's boundary. An infinite network of supply chain processes directly and indirectly influence each life cycle stage but with finite resources, judgements need to be made regarding the system boundary. The problem can be mitigated, to some degree, by adopting consistent criteria across life cycle stages and across different LCAs. However, even when this is done the system boundary cut off for each stage can vary greatly. This can distort the relative significant of overall values attributed to each product or life cycle stage.

If studies make exclusions (as a result of system boundary cut offs) on the basis of data unavailability the usefulness of the overall impact assessment is undermined. More specifically, the exclusion of supply chain processes leads to a "truncation error", meaning that certain emissions are unaccounted for in the final calculated life cycle carbon footprint.

Input-Output Life Cycle Assessment (IO-LCA)

A complementary technique for life cycle analysis is environmentally extended input output analysis (EEIO-LCA). This is considered a "top-down" approach because it uses a macro-economic model to convert economic data from different sectors of the economy into environmental impacts. Used on its own, it suffers from significant generalisations. It assigns a generic financial carbon intensity to goods and services based on their economic sector and country or region of production, but is incapable of differentiating between the carbon intensity of products within each sector and region.

Its strengths are that it does not incur the arbitrary system boundary cut offs inherent in P-LCA, so does not systematically under-estimate. i.e. it provides a system-complete assessment of the upstream supply chain. A further benefit of IO-LCA is that financial data is all that is required.

This approach is useful for assessing complex systems where multiple interdependencies exist, and avoids the risk of data gaps and truncation errors associated with P-LCAs(9). However, IO-LCAs suffer from various limitations. It lacks specificity, applying generalised financial emissions intensity factors to goods and services from different combinations of supplier industry and location(9).

Overall, it is a generic approach, with the advantage of being system-complete because it accounts for all the supply-chain pathways throughout the economy. It is effective at carbon accounting for a product that has a diverse, complex supply chain, and where it is important to produce a complete picture of emissions across the entire supply chain, like in ICT. Although more generic, a thorough IO-LCA does require a significant amount of resources.

Hybrid Life Cycle Assessment

H-LCA seeks to combine the strengths of P-LCA and IO-LCA, counteracting the weaknesses of each on their own. It can either take P-LCA as its start-point and fill in the truncated omissions using IO-LCA or, alternatively, it can take a system complete IO-LCA as its start-point and substitute elements of P-LCA into the model to improve the accuracy and specificity in key areas. The former is often more appropriate for product LCAs whilst the latter often appropriate in company supply chain assessments where the purchase ledger can be used to create an instant, initial, sketch into which bottom up elements can be iteratively substituted. Whichever hybridisation approach is adopted, great care is needed over the system boundary where the approaches are combined to ensure each element is included without double counting. Within any given resource, H-LCA stands to produce a more realistic result than a pure LCA. Hybridisation of a P-LCA does not compromise the specificity of any of its elements, but simply fills in the gaps of 'hard to reach' elements with an estimate drawn from a macro-economic model, rather than discounting them. By delivering system completeness, H-LCA enables a consistent treatment of the system boundary which is essential for comparisons between LCAs. This is especially important when comparing different life cycle stages, or different types of product, since the P-LCA truncation varies greatly between these.

> **H-LCA enables a consistent treatment of the system boundary which is essential for comparisons between**

H-LCAs are an effective tool to understand and address GHG emissions where gaps in data availability required for a P-LCA are significant. They are especially important in the complex and opaque global supply chains that underpin the ICT sector, which amplify the weaknesses inherent in pure P-LCA studies.

Transparency and trust: Why does it matter?

As outlined above, transparency is vital when conducting LCAs. Transparency in the LCA process means ensuring that methodologies, data sources, assumptions, and calculations are openly shared. This openness allows stakeholders, including consumers, businesses, and policymakers, to understand and critically evaluate outputs, fostering informed decision-making. Without transparency, the integrity of the analysis may questioned and meaningful interpretation of the results severely hampered.

One factor that can improve the reliability and comparability of outputs is the implementation of fit-for-purpose standards. These can potentially generate the consistency of metrics and this is essential for achieving comparability across different assessments. There are several standards that aim to increase consistency in LCAs. There are 'foundational' standards (e.g., ISO-14067) which can be applied to any product category. The broadness this necessitates leaves a lot of leeway for LCA practitioners and, therefore, inconsistencies between individual LCAs. Sector-specific standards, such as the GHG Protocol's ICT Sector Guidance, aim to reduce the amount of leeway, however for software LCAs they require further development (for example they contain out of date guidance which is advised against by industry experts (e.g., estimating embodied emissions of equipment based on use-phase emissions)

> **With transparency and consistency, trust in the outputs is created. This is critical as stakeholders must trust the reported data to help ensure carbon reduction actions are communicated and implemented in a consistent manner.**

Accuracy matters – but carbon reduction actions cannot wait.

One of the central debates in selecting a LCA methodology is whether the increased accuracy of a detailed LCA justifies the complexity and cost, or whether a relatively simple analysis provides enough information for effective decision-making to drive decarbonisation.

For cloud providers, and software developers that sell cloud native products, the accuracy of a LCAs can have significant reputational, regulatory, and operational implications. As more organisations adopt sustainable business practices, and investors scrutinise environmental, social, and governance (ESG) performance, software and cloud providers must ensure their sustainability claims are backed by robust data. Focussing on a LCA with a high level of accuracy can be important for reporting and building consumer trust, as well as to set ambitious, science-based climate targets.

However, the additional resource costs required for more accuracy can delay important decisions or carbon reduction initiatives. Furthermore, in rapidly evolving sectors like cloud computing, where hardware is frequently upgraded and operational practices shift, the most accurate data may also quickly become outdated, questioning whether the marginal gains in accuracy are worth the effort.

For software and cloud consumers, the value of LCA accuracy may depend on the organisation's sustainability goals and level of control over its cloud usage. Large consumers of software and cloud services may desire a detailed LCA to understand these purchased goods and services, due to their significance on organisational emissions, to guide decision-making. A company with limited cloud software usage may find it more practical to focus on the sustainability credentials (net zero target and credible transition plan) of their cloud supplier, or a more cursory analysis LCA, i.e., procuring a provider that is more sustainable. The type of LCA method used depends on the goals of the organisation, the availability of data, and the level of control over the software services, and associated infrastructure, used. Regardless of method, the objective remains the same: **identify emission hotspots** to facilitate informed decisions and drive decarbonisation in ways that are both effective and achievable. **Identifying Emission Hotspots: The Path to Carbon Reduction Actions**

> **The true value of LCAs lies in their ability to identify emission hotspots. This knowledge allows both providers and consumers to take targeted carbon reduction actions.**

The data generated by LCAs can help providers pinpoint areas for improvement, such as targeting the supply chain, transitioning to renewable energy sources, or optimising hardware utilisation to reduce idle capacity. However, the software industry cannot rely on cloud providers doing all the work. By conducting LCAs of their own cloud software and usage, they can better understand the environmental impact of the services built and consumed, and take proactive steps to mitigate their impact.

Identifying hotspots also has financial implications. The more is understood about where emissions are concentrated, the more effectively cloud resources can be allocated to support software delivery in a cost-efficient manner. Reducing inefficient use of software and cloud services not only reduces emissions but also reduces operational costs – a win-win for both sustainability and the bottom line.

Furthermore, as regulatory pressures mount, having a granular understanding of the carbon footprint of software and associated cloud operations will be critical. The European Union's Corporate Sustainability Reporting Directive (CSRD), for example, requires companies to disclose detailed sustainability data, including Scope 3 emissions from their supply chains. For organisations that rely heavily on cloud services, failing to accurately account for emissions from their cloud usage could lead to compliance risks and reputational damage.

Thus, accuracy matters – but only up to a point. The goal of a LCA is to inform decisions that lead to meaningful emissions reductions. In some cases, the marginal increase in accuracy associated with P-LCAs may not lead to significantly different actions compared to a well-executed H-LCA, or even an IO-LCA. In this context, the key question is not just about data accuracy, but whether the outputs results in actionable decisions.

A Case Study: A P-LCA and H-LCA approach to measuring the impact of Sage Software

Sage is a software as a service (SaaS) business providing technology and support to millions of small and medium-sized businesses (SMBs) around the world to manage finances, operations, and people. We are market leader for integrated accounting, payroll and payment systems.

As part of our most recent carbon footprint reporting (FY23), we found that our indirect Scope 3 emissions from the Use of Sold Products accounts for 43.9% of our total. As a result, our immediate focus has been to improve the accuracy of product emissions data to aid the development of an appropriate carbon reduction plan for these emissions. The first step in this process was to understand the differences, if any, of the lifecycle carbon impacts between cloud and non-cloud hosted software.

To test the hypothesis that cloud hosted products have lower GHG emissions, we completed a carbon footprint of two software products across their life-cycle stages: Sage 100 FR and Sage Intacct. These two products were selected thanks to the different infrastructure requirements.

Sage 100 (or Sage 100Cloud) is a comprehensive enterprise resource planning solution designed for SMBs. It is adapted to each country, with Sage 100 FR being the version distributed in France and French-speaking countries. It can be hosted either via on-premises or via the cloud.

Sage Intacct is a cloud-based financial management and accounting software designed for SMBs. It is available in five regions but predominantly sold in the US. Intacct is cloud-native and can only be hosted via the cloud.

The study had three objectives:

1. Contribute to the ICT sector's understanding of how to calculate the GHG emissions of software products.
2. Support Sage's future product design and implementation decisions to reduce GHG emissions and support Sage's and its customers' net zero goals.
3. Calculate the emissions of Sage products in order to understand and compare emissions between cloud and non-cloud software and hosting pathways.

To achieve the above objectives in the most constructive manner, Sage decided to complete the carbon footprints by first using a P-LCA method and to later hybridise this alongside various methodological variations. A detailed report outlining the methodology, assumptions, and conclusions of these will be presented in Q1 2025 but a brief outline of each is presented below.

P-LCA approach

The study, completed by Eunomia Research and Consulting, aimed to align with the ISO14067 standard and the GHG Protocol's ICT Sector Guidance. The broad nature of the standards and real-life data availability limitations, meant that bespoke calculations where often required. The study utilised Sage's primary data where available, and relied on secondary research where it was not (particularly for downstream/customer-side impacts). We engaged industry experts at the University of Bristol to validate the inputs and assumptions, and provide peer review of the study.

One of the key findings from the P-LCA was that products operating in the cloud was (as expected) more carbon-friendly than those reliant on on-premise servers. Another key finding was that the results are very sensitive to the assumptions made in the model (and that there is insufficient guidance in existing standards to help with the choice of assumptions), demonstrating the need to replace assumptions with primary data where possible, and to develop specific LCA standards to guide on assumptions where they need to be made.

The process was resource-intensive and if the industry at large are going to generate LCAs for their portfolio, there must be some way of doing this more efficiently. An carbon calculator developed by the industry at large could be one way of doing this.

H-LCA approach

This study was completed by Small World Consulting (SWC) and built on the P-LCA work completed by Eunomia. Its core purpose was to demonstrate the significance of applying different assumptions and the impact of changing the system boundary, including hybridising datasets. The hybridisation applied spend-based factors to supplement the Eunomia outputs, to ensure the outputs were system complete. For example, spending-based factors were used to provide estimated emissions from producing physical devices that re required to operate the software (e.g., docking stations and switches).

One of the key findings from the H-LCA was the importance of accounting for emissions incurred from the date of purchase to 2023, as well as including all parameters associated with developing, producing, and deploying the products. Another was the overall impact of software.

Sage's study identified key recommendations regarding carbon reduction actions for software customers, software providers, and Cloud providers alike.

Conclusion

The ICT industry is uniquely positioned to drive sustainability efforts by embracing carbon lifecycle assessments (LCAs) as a tool to enable **trust and carbon reduction actions**.

It is clear that the cloud infrastructure offers immense potential for innovation, but we must ensure that cloud development strategies are aligned with global efforts to combat climate change. LCAs not only help software and cloud developers, and their users, to understand the full environmental impact of such services, but also empower these stakeholders to make informed decisions about how the software and supporting infrastructure is built and consumed.

Insights gained from LCAs are critical in helping drive meaningful decarbonisation, but the LCA approach, and detailed analysis required, will vary between stakeholders. Detailed approaches such as P-LCAs and H-LCAs (with a significant focus on P-LCAs as part of its methodology) are important for software and hardware manufacturers as they deliver accurate information to the buyer of the product (e.g., server for a datacentre provider or software for a user). However, a H-LCA approach that predominantly focusses on a IO-LCA assessment, may be sufficient for providers and users to stimulate decarbonisation actions.

As a SaaS provider, Sage can address the decarbonisation challenge, in part, by partnering with cloud providers to collaborate on improving the transparency of data provision for Scope 1, 2 and 3 emissions; and/or by taking steps to optimise cloud workloads within their business.

Despite the clear benefits LCAs bring to decarbonisation efforts, transparency with assumptions and data requires further standardisation. Informed, relevant, and trusted carbon reduction actions are only possible with the improved availability of data on cloud emissions. This will steer businesses to choose products and services that have lower emissions. Transparency in data and assumptions is also needed to enable software developers to design more efficient software, and quantify the emissions benefit(s) this brings.

Bibliography

1. Stadig, D., Slijkerman, J.F. & Zhang, C. Data centre growth proves crucial in the shift towards a digital economy. https://think.ing.com/. [Online] ING, 24 October 2023. [Cited: 28 August 2024.] https://think.ing.com/articles/data-centre-growth-proves-crucial-in-shift-to-digital-economy/#a3.
2. Moore's law is the observation that the number of transistors in an integrated circuit (IC) doubles about every two years. https://en.wikipedia.org/wiki/Moore%27s_law. [Online. Cited: 09 October 2024.]
3. Electricity Intensity of Internet Data Transmission Untangling the Estimates. Aslan, J., Mayers, K., Koomey, J.G. and France, C.,. 4, s.l. : Journal of Industrial Ecology, 2017, Vol. 22.
4. The electricity consumption and operational carbon emissions of ICT network operators 2010–2020 and beyond. Malmodin, J., Lundén, D. and Bergmark, P. 1, s.l. : Telecommunications Policy, 2024, Vol. 48, p. p.102538.
5. A review of assessments of the greenhouse gas footprint and abatement potential of information and communication technology. Bieser, J.C.T., Hintemann, R., Hilty, L.M. and Beucker, S.,. s.l. : Environmental Impact Assessment Review, 2023, Vol. 99, p. p.107033.
6. ITU, The World Bank &. Measuring the Emissions & Energy Footprint of the ICT Sector: Implications for Climate Action. Washington D.C. and Geneva : s.n.
7. Freitag, C. et al. The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations. Patterns 2, 100340 (2021).
8. Graham, S., Yashkova, O. & Rutten, P. AI Datacenter Capacity, Energy Consumption, and Carbon Emissions Projections. https://www.idc.com/getdoc.jsp?containerId=prUS52611224#:~:text=IDC%20expects%20the%20surging%20demand,)%20of%2040.5%25%20through%202027. [Online. Cited 09 October 2024).
9. Hybrid life-cycle assessment for robust, best-practice carbon. C. Kennelly, M. Berners-Lee, C.N. Hewitt. 2019, Journal of Cleaner Production, Vol. 208, pp. 35-43.