Open:UK® Open:HQ®

AI Openness: 3 Years on from Llama 2

The Commons we have not (yet) created

Anastasia Stasenko
Founder and CEO, Pleias

Index

Index

2.14.2 Thought Leadership - The Commons we have not (yet) created

Anastasia Stasenko
Founder and CEO, Pleias

The most powerful models Europe lost access to this year happen to be called the Fable series. It is the kind of coincidence you cannot improve on. The suspension did not create a vulnerability so much as expose a fable that had been told for years: that you did not need to build the substrate of artificial intelligence, only to use it well. Own the application layer, the story went, and let others burn the capital underneath. When the layer underneath was switched off from Washington, the story switched off with it.
The lesson most people drew was the obvious one - you cannot rent a substrate and call it sovereignty. That is true, and it is also the least interesting thing to say about it. Of course frontier models are expensive; of course whoever pays for the compute holds the leverage. If the case for decentralised AI stopped at “we need our own GPUs,” it would be a procurement problem, and procurement problems get solved or they don’t. The reason the suspension matters is that it exposed a deeper dependency, one that does not disappear even if you buy the hardware.

Here I have to bring in my philosophical background and talk about the French philosopher Gilbert Simondon, because he described the shape of the problem before there was a machine to attach it to. Individuation - the process by which a subject actually forms - is, in his account, two things at once. It is collective: a mind does not form in isolation but through feedback loops with a milieu, other minds, other practices, the slow friction of people who do not agree. And it is contextual: it happens inside a language, a craft, a tradition dense enough with shared references to sustain real disagreement. There is no individuation in general. There is only individuation within a milieu specific enough for thought to occur.

Hold those two commitments against the frontier proposition - one model, trained on everything, served to everyone through a single interface - and the mismatch is structural, not economic. A model like that can be extraordinarily responsive and still be uncontestable. You can prompt it; you cannot argue with it the way you argue with a colleague or a tradition. It flattens the milieu into a generic surface. Close the loop through which a subject individuates inside one company’s API, and the individuation that results is not free. It is administered.

This is the argument for decentralisation that does not reduce cost or anti-monopoly politics. It is a claim about the conditions under which people think. You need models embedded in particular languages, particular domains, particular communities of practice - not because the local is romantically superior to the universal, but because the loop through which a subject forms requires a milieu specific enough to sustain genuine collective thinking. One model trained on everything by a single entity cannot provide that. It can simulate responsiveness. It cannot simulate the situated friction of a real epistemic community. The frontier logic is not only concentrated. It is epistemologically impoverished - and plurality is not a concession to scarcity but a condition of the kind of thinking that matters.

The strange thing, the thing that gives this argument political weight it would otherwise lack, is that the CFO arrives where the philosopher does. As AI stops answering questions and starts running operations - procurement, compliance, logistics, decisions at every level - the model a company operates through becomes a layer of its governance. Every inference call constrains a process. No firm should run its core operations through a provider that reads its data and learns its patterns while serving its competitors on the same infrastructure. The commons argument says closed cognitive infrastructure is incompatible with collective thought. The competitive argument says closed operational infrastructure is incompatible with autonomy. Simondon and the CFO, by entirely different routes, demand the same substrate: plural, contextual, governed by those who depend on it.

OpenUK AI Openness: 3 Years on From Llama 2, July 2026

The uncomfortable part is that wanting this does not make it buildable, and the tradition we reach for first does not fit. We call the alternative open source, but there is no open source AI in the sense there is open source software. Every capable open-weight model - Llama, Granite, the Mistral weights, Qwen, and OLMo, the one that comes from a research institute rather than a company - is a corporate release with a permissive licence attached, and the licence is the whole fragility. Open source software worked because the people who built it owned the means of producing it - a laptop, a connection, time - and because the GPL made openness irreversible: once code was free it could not be enclosed again. There is no GPL for weights. You can open them, but if the data, the methodology, and the training code stay closed, what you have is an artifact a corporation can withdraw tomorrow. A community built on a revocable gift is not a commons. It is a tenancy.

So the commons tradition has to borrow from one it has always ignored - the tradition that knew how to provide things too expensive for any individual and too important to leave to monopoly. Railways, electricity, telecommunications each posed exactly this, and each was answered with an institutional invention: common carriage, public utilities, cooperatives, rural electrification. Capital-intensive infrastructure that is a precondition for participation cannot be governed by the same logic as the activities it enables. It is worth noting who else has reached this conclusion. When Leo XIV placed AI infrastructure, data, and algorithms under the universal destination of goods, a tradition working from entirely different premises landed exactly here. Compute has a universal destination, or it has an owner.

And the challenge is double, because it is not only compute. It is also data - the permissibly licensed material models learn from, before it is scraped into proprietary pipelines and sold back as a service with nothing returning to the commons that produced it. Common Corpus, the two-trillion-token open dataset we built at Pleias, is one very modest attempt at that second layer: a demonstration that what models learn from can itself be open, and governed by something other than extraction. But without public compute, even open data feeds an enclosed pipeline. The two commons rise or fall together.

What is missing is the institutional form, and it is within reach precisely because the software world already invented most of it. A copyleft licence for AI: use the weights, modify them, build on them, but derivatives stay open, and commercial serving returns a share to the commons that produced them. Couple that to cooperative governance and an inference-revenue loop - training as the capital cost, inference as the demand that funds the next run - and the commons sustains itself rather than begging for subsidy, the way energy cooperatives and credit unions always have: those who depend on the infrastructure govern it, and the surplus stays.

None of this exists at full scale. All of it is being attempted in fragments. Whether those fragments converge into a form with the weight to hold is a political question, not a technical one - and for the first time the coalition is real, because the philosopher who wants plural infrastructure and the company that wants operational autonomy have finally begun to ask for the same thing.

6. Formalities

6.1 Authors of this Report

Professor Amanda Brock, CEO, OpenUK

OpenUK CEO, Amanda's built one of open source's most recognised and impactful organisations. Executive Producer of State of Open Con (2023- 2025), Amanda's a globally sought-after keynote speaker.A lawyer with 25 years' experience, 5 as GC of Canonical, she's been instrumental in shaping open source's legal frameworks, as she was internet law during the early 2000's. Regularly contributing to tech press, she edited 'Open Source: Law, Policy and Practice', (2022).

Recognition: Computer Weekly 50 Most Influential Women in UK Tech (2023, 2024); Computing IT Leaders 100 (2023, 2024); Lifetime Achievement Award WIPL (2022); Women Who Will Changemaker (2023); INvolve Heroes (2022, 2023); Novi Awards (2024) and Ambassador, Open Charge Alliance. Advisory Appointments: UK Cabinet Office Open Standards Board; UKRI Digital Research Infrastructure; UKRI Exascale; KDE; commercial boards – Mimoto, Scarf, FerretDB and Space Aye; and is Fellow Open Forum Academy; Distinguished Fellow Rust Foundation; and European Representative, OIN.

Dr. Jennifer Barth, Head of Research, Symmetry Research at FSP and Chief Research Officer, OpenUK

Jenn holds a DPhil from the University of Oxford and is Head of Research at Symmetry Research at FSP, where she brings a rare blend of ethnographic depth, commercial acuity, and ethical clarity to the AI conver- sation. Jenn's work sits at the intersection of technology, organisational culture, and socioeconomic change. She is known for translating complex technological shifts into human-centred insights that resonate deeply with executive audiences. Her measured, critically informed perspective — grounded in qualitative research and lived experience — enables leaders to move from abstract concern to practical, principled action. Alongside her role at FSP, Jenn serves as Chief Research Officer at OpenUK.

Karan Saini, Research Manager, OpenHQ

Karan Saini is a New Delhi-based security researcher, public interest technologist, and open source contributor. He was most recently a Senior Information Controls Fellow with the Open Technology Fund, hosted by the Internet Governance Project at Georgia Tech, where he researched the scale of DNS-based web censorship in India. He regularly writes for leading publications and has contributed to award-winning media projects.

6.2 Creators of this Report

OpenUK

OpenUK is the unique open tech industry organisation for the business of open technology in the UK. It spans the opens – software, hardware, data, standards and AI and is the convening point for the UK's business, academic and contributing communities across open tech. Our work supports the UK's journey to become "The State of Open". Our organisation is run with the support of our volunteer community and their leadership in the tradition of open source manner delivering on three pillars: community, legal and policy and learning. OpenUK's Community is recognised through our world-leading recognition program including the Open UK Awards (the Oscars of Open Source) now in their 7th year, New Year's Honours Lists and Ambassador Scheme.
OpenUK undertakes research and reporting both on its own account through its "State of Open Reports" and on a commissioned basis for third parties. Case studies, Thought Leadership, Surveys and desk-based research are included in our reporting which pushes the envelope and leads the way. Our Research and Reporting Show and Tell events coalesce the global open source research communities digitally to regularly update and share research practices and topics. .

The State of Open Con has become one of the world's leading open source conferences since its inception by OpenUK in 2023. In 2027 we expect to host 1000 people across 8 tracks and plenary sessions, with at least 50 partners in our delegate experience space and over 200 speakers in London. For 2026 we are on the road, with smaller events across the UK, meeting our audience. Contact mailto:admin@openuk.uk

Symmetry Research

Symmetry Research, an FSP company, looks beyond the surface and behind the curtain of the fundamental innovations and trends shaping our society, markets, culture, and values. We are academics and researchers looking at the intersections of emerging technology and socioeconomic impact, producing independent research for thought leadership and business solutions. We take a research-led approach to helping organisations transform through and with technology. Symmetry Research's mission is to share and grow knowledge about the interaction of technology and everyday lives. We want to understand the past, present, and future of human interaction with emerging technologies and socioeconomic changes-from behaviour to context, nature to nurture, origin to experiences-we do independent research for high-stakes business decisions.

6.3 Contributors

Anastiasia Stasenko
Anastasia Stasenko is the co-founder and CEO of Pleias, a Franco-German startup developing full open source small reasoning models trained exclusively on permissively licensed data. A senior associate lecturer in data analysis and digital strategy at Sorbonne-Nouvelle University, she coordinated the creation of Common Corpus, the largest fully open dataset for training LLMs, with over two trillion tokens.

Professor Anil Madhavapeddy
Anil Madhavapeddy is the Professor of Planetary Computing at the Department of Computer Science and Technology in the University of Cambridge.

David Soria Parra
David Soria Parra is a Member of Technical Staff at Anthropic based in London. He co-created the Model Context Protocol (MCP). He leads its technical direction, governance, and ecosystem development. Before joining Anthropic in 2024, David spent nearly a decade at Meta leading engineering organizations, and served as an Entrepreneur in Residence at Sutter Hill Ventures. He is a long-time open source contributor with a track record spanning PHP, Git, and Mercurial.

Hannah Foxwell
With over a decade of experience in engineering leadership, Hannah Foxwell has always advocated for the human aspects of technology transformation and evolution.Hannah is relentlessly curious about the tools, technologies, processes, and practices that make life better for the people involved in software development. This curiosity has naturally evolved alongside technological changes—from Agile to DevOps, DevSecOps, Cloud Native, and Platform Engineering. In 2026 Hannah co-founded BIMP, bringing together her knowledge of cybersecurity, platform engineering, and AI to deliver software supply chain security at the speed of AI.As the creator of "AI for the rest of us" Hannah leads a community of over 2,000 enthusiasts in the UK, providing free monthly events and an annual conference aimed at demystifying the technology and improving AI fluency.

James Drayson
James Drayson is the CEO and co-founder of Locai Labs, a UK company based in London that develops on-prem sovereign AI so enterprises can own their intelligence. James holds a Bachelor of Science in Biology from the University of California, Los Angeles, and a Master of Science in Innovation, Entrepreneurship and Management from Imperial College London. He is Forbes featured and a UKRI Young Innovator Award winner.

Matthew Barker
Matt Barker is co-founder and CEO of BoltMCP, focused on helping enterprises securely connect AI to their data and control how it's used. Previously, he was CEO and co-founder of Jetstack, where he helped build one of the most widely adopted projects in the cloud native ecosystem, cert-manager. Under his leadership, Jetstack bootstrapped to a team of 50, and was acquired by Venafi in 2020.

Rohit Ghumare
Creator & Maintainer, Agentmemory, CreateMVP, and AI Engineering from Scratch AI-native course. Open source Leader and Cloud-native engineer. He writes and talks about AI agents, harnesses, AI Inference, and the infrastructure underneath, and is currently building iii.dev. Contributor to major open source global projects.

OpenUK AI Openness: 3 Years on From Llama 2, July 2026

Sachiko Muto
Sachiko Muto is the Chair of OpenForum Europe, a senior researcher at RISE Research Institutes of Sweden, and a Board member of the Drupal Association. She originally joined OFE in 2007 and served for several years as Director with responsibility for government relations and then as CEO. Sachiko has degrees in Political Science from the University of Toronto and the London School of Economics; she received her doctorate in standardisation policy from TU Delft.

Sal Kimmich
Sal's work is centered on advancing secure computing practices and pioneering privacy-enhancing technologies to safeguard sensitive data. At the Confidential Computing Consortium (CCC) within the Linux Foundation, they play a pivotal role in the development and implementation of Open Source projects that establish Trusted Execution Environments (TEEs). These environments are designed to provide secure areas within hardware that protect data in use, addressing critical security challenges across various industries.

Vincent Koc
Vincent Koc is Chief Architect at the OpenClaw Foundation, where he leads engineering for one of the fastest-growing open source projects in AI. An engineer and technologist with nearly two decades of experience across the Asia Pacific and the US, he has contributed to major open source projects across the data and AI ecosystem and teaches applied AI at universities including MIT. He joined OpenClaw as a community maintainer focused on security and observability before moving to the Foundation full-time.

Quoted Individuals
Andrew Martin, Co-Founder and CEO, ControlPlane; Adrian Mouat, DevRel Engineer, Chainguard; Chris Yiu, Head of Policy Northern Europe, Meta; Eric Schmidt, Former CEO of Google, Chair of the Special Competitive Studies Project; Jean-Baptiste Kempf, President, VideoLAN and Kyber; Jensen Huang, CEO, NVIDIA; Juliette Denny, Founder, Growth Engineering and Zavmo; Luke Hinds, Founder, nolabs, Inc; Sam Altman, Founder and CEO, OpenAI; Toran Bruce Richards, Founder and CEO, AutoGPT; Yann LeCun, Founder and Chairman, AMI Labs

6.7 Sponsors

OpenUK is grateful to its many sponsors and supporters without whose support its work including its report would not have been possible in particular, its general sponsors Arm, GitHub, Google, Linaro, Microsoft, Red Hat.

Download the full report here

openuk.uk/report/aiopenness2026

OpenUK Newsletter

Our Conferences - State of Open Con

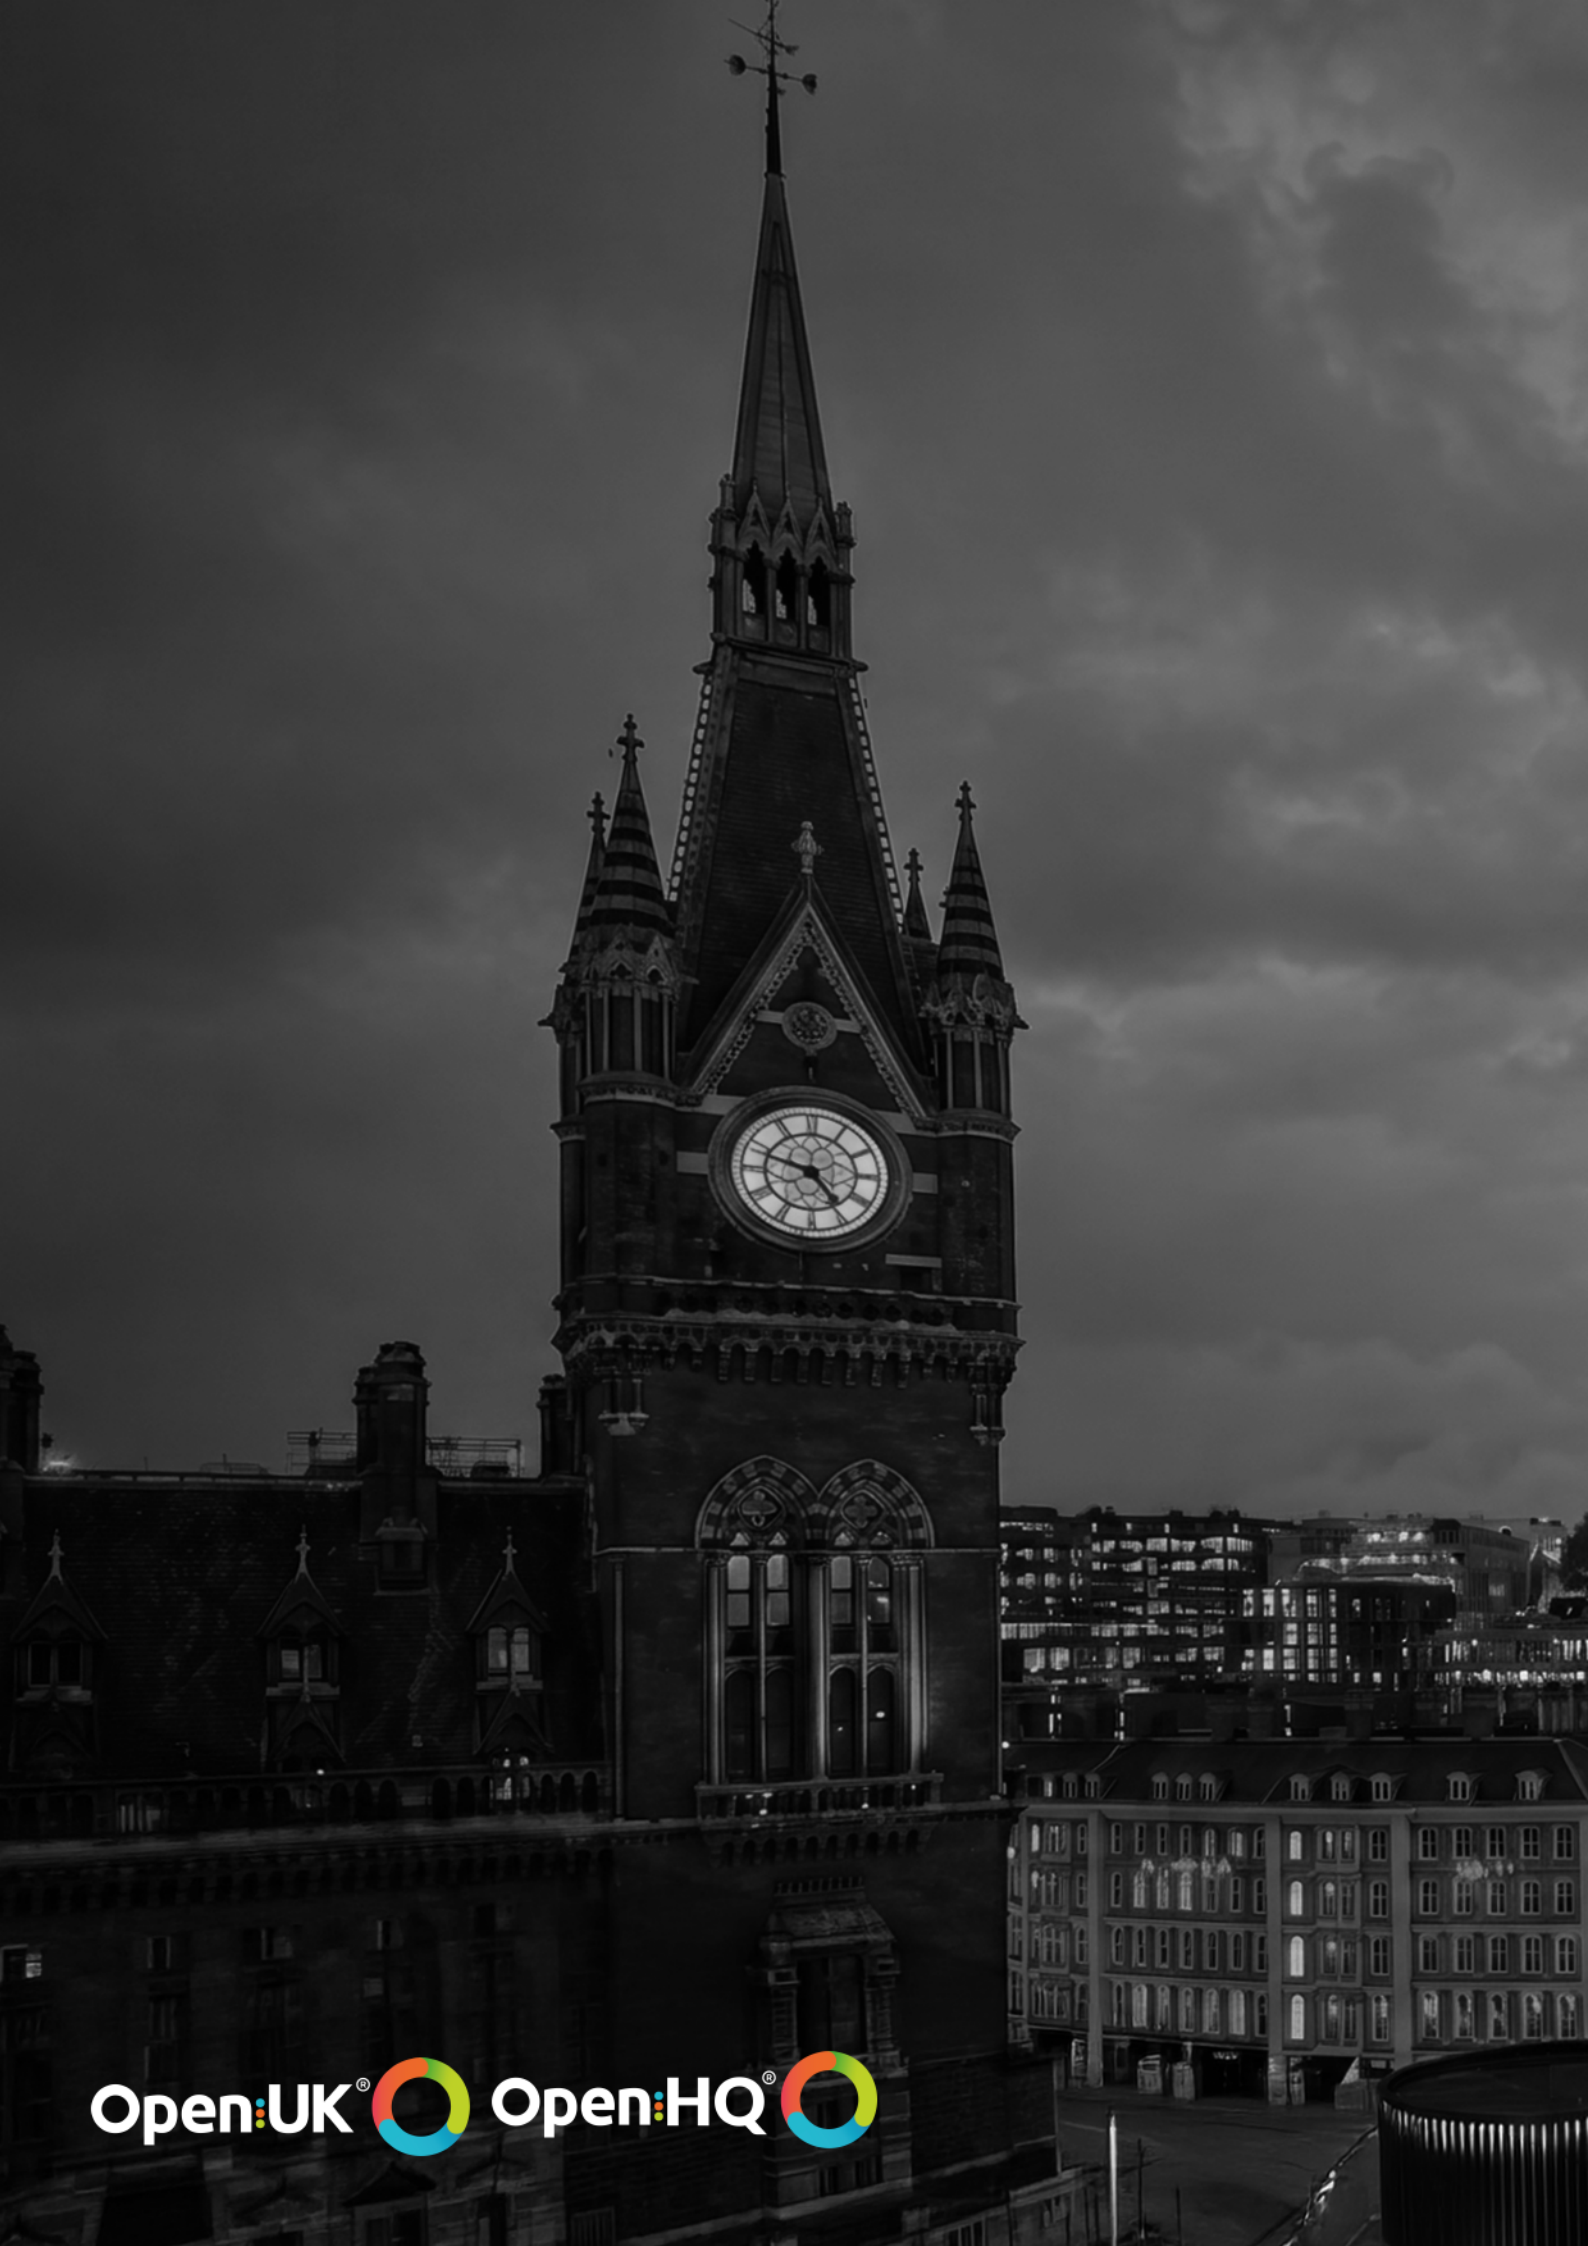
Open:UK®
Open:HQ®