



Department for
Science, Innovation
& Technology



UK Research
and Innovation

UK Compute Roadmap

CP 1352



UK Compute Roadmap

Presented to Parliament by the Secretary of State for Science, Innovation and Technology by
Command of His Majesty

July 2025



© Crown copyright 2025

This publication is licensed under the terms of the Open Government Licence v3.0 except where otherwise stated. To view this licence, visit nationalarchives.gov.uk/doc/open-government-licence/version/3.

Where we have identified any third party copyright information you will need to obtain permission from the copyright holders concerned.

This publication is available at www.gov.uk/official-documents.

Any enquiries regarding this publication should be sent to us at compute@dsit.gov.uk

ISBN 978-1-5286-5894-2

E03400533 07/25

Printed on paper containing 40% recycled fibre content minimum

Printed in the UK by HH Associates Ltd. on behalf of the Controller of His Majesty's Stationery Office

Contents

Contents	3
Foreword by the Secretary of State for Science, Innovation and Technology	4
Executive Summary	5
1. Introduction	7
1.1 A global race	7
1.2 Ensuring the UK is a leader, not a laggard	8
2. Building a modern public compute ecosystem	9
2.1 Our starting position: a fragmented system in need of direction and rejuvenation	9
2.2 The vision: building a platform for science, innovation and discovery	9
2.3 The UK public research compute ecosystem of the future	16
3. Putting compute to work – powering innovation across the public and private sector	17
3.1 The case for change – moving to a more strategic and outcomes focused approach	17
3.2 Allocation priorities	18
4. Building AI infrastructure to keep the UK at the cutting edge of AI development	20
4.1 Global and technological context	20
4.2 Our approach – moving now to build future proofed capacity	20
4.3 Going beyond infrastructure to deliver widespread national and local benefits	22
5. Creating sovereign, secure and sustainable capability	24
5.1 Our approach to sovereignty – pragmatic, balanced and strategic	24
5.2 From prototype to deployment – a pull through mechanism for UK suppliers	25
5.3 Compute – An opportunity area for Sovereign AI	26
6. Conclusion and next steps	28
7. Compute Roadmap – at a glance	29

Foreword by the Secretary of State for Science, Innovation and Technology

Artificial intelligence is reshaping the world around us – transforming the way we innovate, work, and solve some of the greatest challenges of our time. From accelerating the development of new medicines to supporting more efficient public services, the potential of AI to drive economic growth, improve lives, and secure our national resilience is vast.

But realising that potential depends on the infrastructure we build to support it. Compute – the processing power behind every AI breakthrough – is not simply a technical resource. It is a critical enabler of scientific excellence, industrial competitiveness, and sovereign capability. Countries that can deliver and direct world-class compute infrastructure will shape the future of AI, not just consume it.

This roadmap sets out how the UK will rise to that challenge. It is a bold, long-term plan to transform our national compute ecosystem – upgrading public research capacity for AI and wider science, unlocking new investment through AI Growth Zones and ensuring that the UK is home to next wave of leading AI and compute companies.

It also plays a key role in enabling the government's broader plan for change – driving innovation, growth and opportunity across the economy. By providing the foundations for AI adoption across science, the NHS, education, defence, and more, this roadmap supports our vision for a more innovative, productive, and resilient United Kingdom.

We are already taking action. With up to £2 billion of investment over the coming years, we are building momentum – from launching the AI Research Resource to revitalising communities across the UK through new AI Growth Zones. We will act with pace and purpose – but also with a clear commitment to security, sustainability, and public value.

This plan will put compute to work for Britain – enabling the science, services, and industries that will define our future and secure the UK's leadership in the age of AI.

The Rt Hon Peter Kyle MP

Secretary of State for Science, Innovation and Technology



Executive Summary

The UK's AI Opportunities Action Plan presented a clear statement of intent to lead in the age of AI. Since it was published, there have been significant geopolitical, economic and technological developments – all of which heighten the importance of developing and deploying AI capability.

Central to this is the need for computing infrastructure that is fit for purpose, not just for today's needs, but for the scale, speed and opportunities of the AI-enabled economy of the future. This roadmap sets out the government's plan to put this infrastructure in place.

Our vision is to build a world-class compute ecosystem that provides a platform for innovation, growth and opportunity across the economy. It will be outcomes-driven with compute directed toward our highest priorities and most transformative opportunities. Our compute ecosystem will support a broad range of users, at the frontier of AI research as well as those adopting AI into their daily workflows. It will have a diverse and resilient mix of compute infrastructure services – spanning public and private systems, AI training as well as inference, national platforms and regional innovation hubs. Crucially, it will be dynamic and adaptable, evolving to support the technologies of tomorrow.

The government has committed up to £2 billion between now and 2030 to build a modern public compute ecosystem, building on the delivery of two new AI supercomputers that form the first phase of the AI Research Resource. We have already started the process to identify and deliver AI Growth Zones across the UK, building on the £44bn of private sector investment in AI Data Centres over the last 12 months. We are also addressing the energy needs of AI through the AI Energy Council and committing to a new generation of advanced nuclear power projects.

To build on this momentum and ensure our ambition matches the rate of scientific and technological progress, we are delivering a 10-point plan that sets out our ambition and provides clarity and certainty to researchers, industry and investors alike.

Building a modern public compute ecosystem

1. **Invest up to £2 billion to deliver a diverse, joined-up and user centred compute ecosystem** – including over £1 billion to expand the AI Research Resource (AIRR) 20x by 2030 and up to £750 million for a new national supercomputer service in Edinburgh.
2. **Establish National Supercomputing Centres** to curate datasets, build software assets and deliver a skills pipeline to support a broad range of users to access and utilise computing power.
3. **Partner with like-minded countries and computing centres** to expand access to a broader range of research infrastructure, facilitating collaboration and skills and knowledge exchange.

Putting compute to use, powering innovation across the economy

4. **Introduce a refreshed allocation model that targets compute toward the UK's highest-impact research and innovation priorities**, backing mission-driven research and strategic national goals.
5. **Guarantee dedicated compute access for the Sovereign AI Unit and the AI Security Institute**, ensuring the UK's core AI functions have the capacity they need to lead.

Building AI infrastructure to keep the UK at the cutting edge of AI development

6. **Deliver large scale AI Infrastructure via AI Growth Zones (AIGZs) across the UK**, ensuring the UK has the capability to support AI training and inference workloads.
7. **Explore new models for delivering the energy infrastructure that powers AI**, including through sustainable solutions such as renewables, advanced nuclear, and innovative grid solutions.
8. **Harness AIGZs to deliver both national and local benefits across the economy**, serving as platforms for innovation, adoption and collaboration through which AI capabilities are developed, trialled, deployed and scaled.

Creating sovereign, secure and sustainable capability

9. **Support British companies to develop sustainable and secure compute technologies**, using the AIRR and AI Growth Zones to provide access, testbeds, and opportunities to scale.
10. **Establish compute as a priority area for the UK Sovereign AI Unit**, supporting research and innovation in new computing paradigms to help British companies grow into global leaders.

Through these 10 actions, we will ensure that the UK's infrastructure allows our scientists and innovators to thrive – and that the next great inventions, AI breakthroughs, and transformative ideas come from the United Kingdom.

1. Introduction

Compute is the engine of the modern digital world. It refers to the processing power delivered through chips, stored in data centres and accessed by software on our smartphones, laptops and by industrial devices. It powers our everyday activities, navigating journeys, connecting with others, and accessing vital public services, as well as the research that improves our lives, such as understanding climate change and discovering new drugs.

The UK is a major global consumer of computing services, with Greater London ranked among the world's top five data centre markets and a long history of research leadership in the application of compute to science.

Compute has emerged as one of the most critical inputs for the development and deployment of advanced AI. Major leaps in capability have been enabled by access to vast computational power. As AI models grow more complex and more widely integrated into the economy, demand for compute is set to accelerate even further. This is not a passing trend, demand for compute at the frontier of AI is set to increase 10,000 times by the end of the decade. At the same time, compute is becoming essential to modern science from genomics to climate modelling to materials discovery, and the science of AI itself.

The government has also set out a Modern Industrial Strategy, a Ten-Year Infrastructure Strategy, and sector strategies covering health, defence and more. Compute, and the applications it enables, will be integral to delivering on all of these plans, because across AI, healthcare, transport, defence, finance, education and in use cases we cannot yet anticipate, **the direction of travel is clear: more users, more and bigger use cases – more demand for more compute.**

1.1 A global race

Nations are beginning to view compute as a strategic enabler and are moving quickly. The USA, Canada, and the United Arab Emirates have each committed to multibillion dollar investments in public compute capacity. Alongside this, hyperscalers have committed over \$300 billion in 2025 alone to build compute infrastructure.

This scale of investment is reshaping global supply chains – from semiconductors and advanced materials to critical minerals, energy infrastructure and data centre technologies – and bringing about a new class of global consortia: alliances between hyperscalers and AI labs, partnering with chipmakers and energy firms, and funded by the largest private and sovereign capital institutions. Entire sectors are being reconfigured to serve the world's growing need for compute as global demand surges.

1.2 Ensuring the UK is a leader, not a laggard

In this context, the role of government is more important than ever. The UK is home to world-class researchers, a dynamic AI ecosystem, and emerging strengths in AI security, cyber and in AI chip design. But without the right infrastructure backbone, these assets will not reach their full potential.

Building that backbone will take time and is only possible with strategic coordination that drives concerted action. The UK's compute ecosystem must evolve – from fragmented and uncoordinated to expansive and mission-driven, serving our public services, researchers and industry.

This roadmap lays out a plan to deliver a national platform for scientific leadership, long-term growth and strategic resilience. It is made up of four purposeful steps:

- **Building a modern public compute ecosystem** providing the capacity and certainty UK researchers and innovators need;
- **Putting compute to use** ensuring we shape science to be more ambitious and impactful in areas of most strategic value and importance;
- **Building cutting-edge AI infrastructure** so that the UK can play its role at the frontier of AI and enable the rapid adoption of AI across the economy;
- **Creating sovereign, secure and sustainable capability** to shape the future of AI and compute, backing the next great British companies.

2. Building a modern public compute ecosystem

Backed by up to £2 billion, we will deliver a modern public research compute ecosystem providing a diverse range of resources to support the full research lifecycle, including compute, data, software and skills. We will deliver a stable and resilient research platform, through a pipeline of investments that provide confidence, certainty and the capabilities that researchers need.

2.1 Our starting position: a fragmented system in need of direction and rejuvenation

The UK's public research compute ecosystem is a critical national asset. Yet while the current landscape has pockets of excellence, the system is strained and incapable of supporting our ambitious research agenda.

Previous reviews of the UK's public compute landscape and technological trends identified several factors that are holding the UK ecosystem back:

- **A lack of capacity and scale** – today's ecosystem lacks the power and bandwidth to serve our researchers. Systems are routinely oversubscribed resulting in researchers looking elsewhere to support their needs or choosing to limit their ambition.
- **Fragmentation and silos** – the landscape is split between discipline-specific systems that operate in isolation. This limits efficiency, hinders collaboration, and prevents the UK from delivering on cross-cutting, mission-driven science.
- **Barriers to access and transparency** – many users, especially from industry or interdisciplinary fields, struggle to navigate the system. Access is unclear, inconsistent, and often closed to new entrants.
- **Untapped strength in software, data and skills** – the UK has a strong base in research software engineering, public data, and AI talent. But these assets are not well-integrated or utilised – missing an opportunity to utilise high value assets and supercharge innovation.

In addition, our public research compute infrastructure is also facing a cliff-edge, with the majority of systems reaching end of life over the next 18 months.

2.2 The vision: building a platform for science, innovation and discovery

We will move from a fragmented and ageing system to one that is confident, collaborative and catalytic. It will be built around five core principles:

- **Stable, long-term funding** – we will establish a multi-year funding model, preventing cliff-edges in the future. Regular refresh cycles will ensure national systems and supporting capabilities stay up to date, with sufficient lead in time for planning and delivery.
- **A diverse and joined-up ecosystem** – we will build an ecosystem that serves the needs of the full scientific and innovation lifecycle – from early-stage experimentation, to testing and development to deployment. Systems will be designed to work together as a coherent, flexible suite, accessible across disciplines and sectors.
- **User centred support and expertise** – we will expand national capacity in user support through a focus on excellence in software engineering and by building a talent pipeline that ensures researchers and innovators can fully leverage the capacity made available to them.
- **Data driven** – we will maximise the value of the UK's data assets by integrating compute with access to large-scale, high-quality datasets. We will also use compute to generate new datasets through simulation, modelling, and AI-enabled research, creating powerful resources for science, innovation, and public good.
- **Dynamic and agile** – we will respond to and incubate emerging technologies, including quantum computing and novel computing paradigms, and provide the space and tools for researchers and developers to experiment with new solutions and apply them to enhance their research.

Following these principles, we will establish a single national compute ecosystem defined by a broad, powerful mix of compute services that are fully integrated. This ecosystem – through centres of excellence, user support services, and programmes for data, software, and skills – will empower even complete novices to use AI and large-scale compute to solve problems, pursue ideas, and make the next great discoveries.

This transformation is already underway. Through the delivery of the AI Research Resource the government is providing access to thousands of new AI chips, many of which have been in use over the previous six months through a more agile and user centric early access model. This Government has launched new programmes to train research technology professionals and made investments in new datasets – such as OpenBind – to catalyse research. However, much more action is needed.

Action 1 – Invest up to £2 billion to deliver a diverse, joined-up, and user centred compute ecosystem – including over £1 billion to expand the AI Research Resource 20x by 2030 and up to £750m for a new national supercomputer at Edinburgh.

We will invest over £1 billion to expand the AI Research Resource (AIRR) twenty-fold – increasing the AIRR's computing capacity increase from 21 AI ExaFLOPS in 2025 to 420 AI ExaFLOPS by 2030.

AI ExaFLOPs Explained:

An exaFLOP is a measure of computing power, representing one quintillion (a 1 followed by 18 zeros) calculations per second.

If one billion people (roughly the combined population of the UK, EU, US and Japan) each had a calculator and performed one calculation per second – then it would take them 665.45 years to do what the current AIRR could do in one second with its power of 21 AI exaFLOPs.

Once the AIRR capacity rises to 420 AI exaFLOP, it would take one billion people 13,316 years to do what the full AIRR will do in one second. That means all one billion people would have needed to start calculating more than 8,000 years before Stonehenge was built – without taking a break.

Procurement for the AIRR will follow a blended model, combining both purpose-built AI supercomputers and cloud-based compute to deliver flexibility, performance, and resilience. By combining both, we will build a resilient, responsive compute ecosystem – capable of adapting to evolving needs while maintaining strategic direction and public value.

Dedicated AI supercomputers will continue at the core of our national capability – providing the high-performance, tightly coupled infrastructure needed for AI model training and mission-critical workloads. At the same time, cloud solutions will play an essential complementary role, offering flexibility, burst capacity, and accessibility for a wider range of users, particularly for inference and more distributed workloads.

To deliver this expanded capacity, we will aim to leverage the potential of public-private partnerships, recognising growing opportunities to partner with private sector providers to design, build, and potentially operate compute systems, while still ensuring public access, strategic control, and security. Such partnerships will maximise the impact of our public investment, bring both technical and commercial expertise into the delivery and operation of national assets, while also creating growing UK-based capability across the compute ecosystem.

How we invest – and in which types of systems – will evolve over time, responding dynamically to market and technological trends. In making our choices, we will be led by the following principles to deliver a best-in-class, balanced, and workload-optimised compute portfolio – designed to deliver value, flexibility, and strategic impact across the UK's AI and research ecosystem:

- **Diversity of suppliers** – we will work with a broad range of vendors and hardware providers to avoid over-reliance on any single supplier, encourage competition, and ensure long-term resilience and innovation.
- **Balance between training and inference** – the AIRR will maintain a flexible mix of systems optimised for both AI training and inference.

- **Support for specialised and novel solutions** – a portion of AIRR capacity will be reserved for emerging, high-potential technologies – including novel chips, architectures, and software stacks – helping to incubate next-generation compute solutions. This will enable us to get ready for the greater role we envisage emerging technologies to play in the UK's compute mix in the future, such as quantum computing.
- **Blending power and purpose** – we will invest in a mix of technologies, including the latest, most powerful chips for large-scale AI training, alongside more efficient or specialised architectures optimised for other lower-intensity tasks. This balanced approach will ensure we meet a wide range of user needs while maximising value for money, deploying the right hardware for the right job.
- **Security and sustainability by design** – as the AIRR grows, all systems will be developed with trusted, cyber secure environments and clear sustainability metrics, reflecting the UK's values and strategic goals.
- **Interoperability and portability** – systems will be designed with a focus on user flexibility and interoperability, ensuring that workloads can move easily between on-premise, cloud, and hybrid environments.

Investment will scale up over this spending review period and is expected to reach over £300 million in 2029. To kickstart the next phase of delivery, we will:

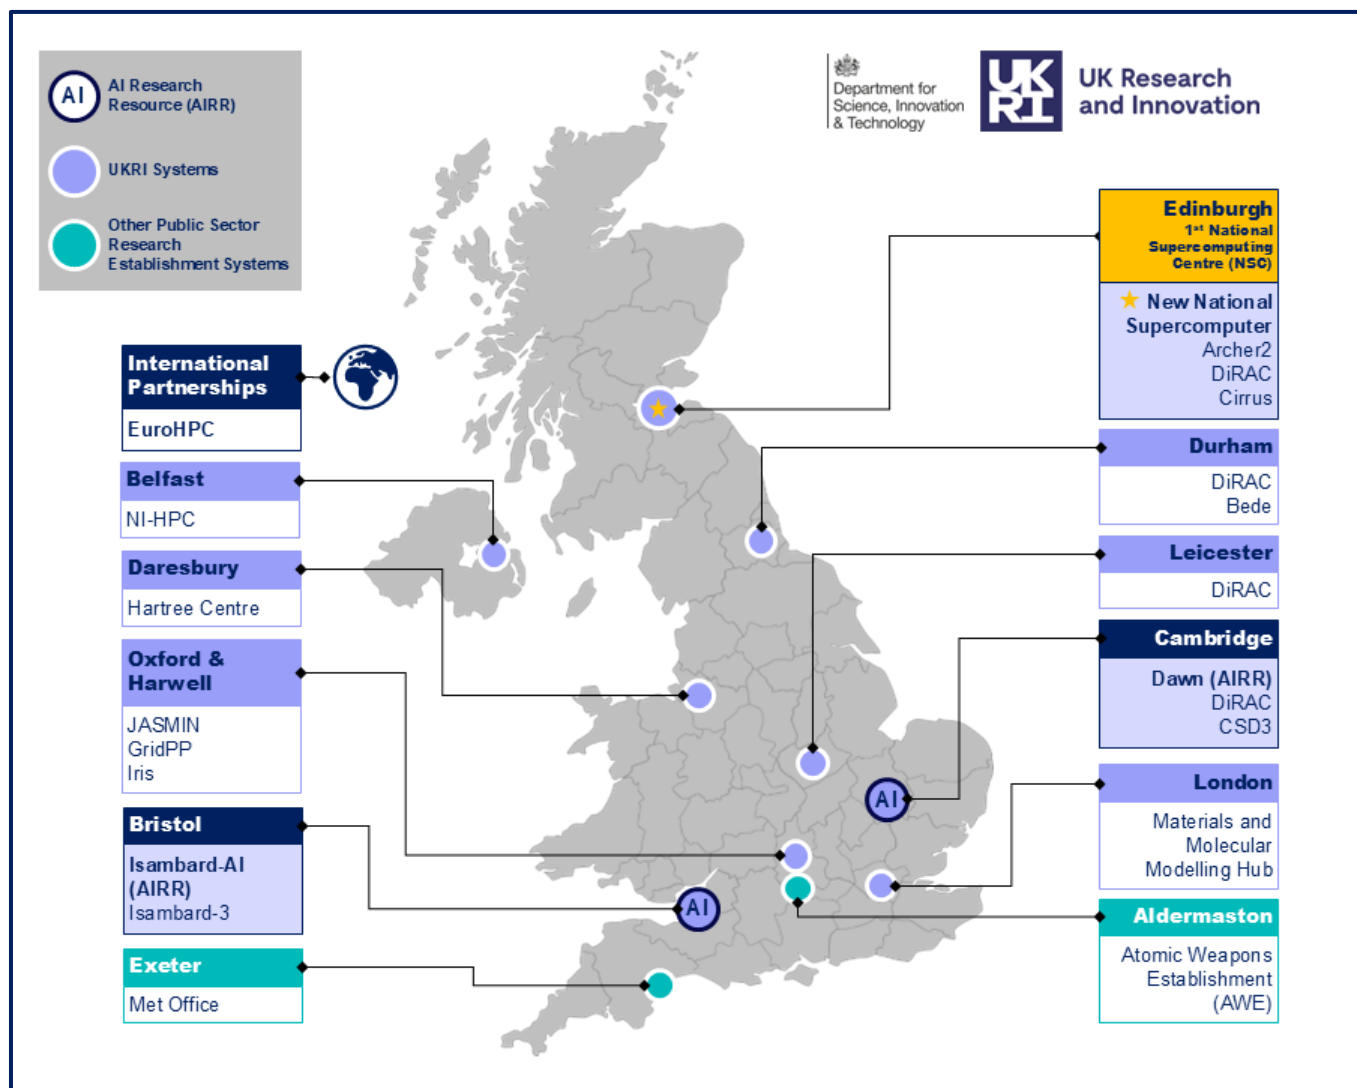
- **Launch a programme to identify opportunities for Public-Private Partnerships** that could coordinate the delivery and operation of new AIRR supercomputing capacity, enabling build out at pace and with agility;
- **Commence procurement of cloud compute to complement our dedicated AI supercomputers** ensuring the AIRR offers both high-performance, purpose-built infrastructure and flexible, on-demand resources to meet the full range of user needs.
- **Deliver 'AIRRPORT'**, which will provide a single front door to the AIRR service. AIRRPORT will provide a seamless experience for users, enabling them to submit jobs, move workloads, access tools, and manage data across AIRR systems ensuring accessibility and interoperability.

The Spending Review committed up to £750 million into a new national supercomputer service at the Edinburgh Parallel Computing Centre (EPCC). The new national supercomputer will replace the current ARCHER2 supercomputer and ensure the UK scientific community retains access to the most advanced modelling and simulation capabilities. This new system will provide sufficient capacity to enable world-leading research in both depth and scale across the scientific landscape and be available for industrial access. The system is planned to come online in 2027 and utilise the next generation of advanced chips that can service both traditional modelling and simulation, and AI workloads – with a projected capacity that would put it in the top 5 supercomputers of today.

Alongside the new national supercomputer, DSIT and UKRI will continue to support a diverse portfolio of public sector research establishment (PSRE) and smaller-scale systems. This will ensure that the UK's public research compute ecosystem retains the range of shared national capabilities needed to service the diverse user needs of different scientific

disciplines – and that flagship PSRE systems such as the Met Office supercomputer are joined-up with the wider ecosystem.

Figure 1: The UK Public Compute Ecosystem



Action 2 – Establish National Supercomputing Centres to curate datasets, build software assets and deliver a skills pipeline to support a broad range of users to access and utilise computing power.

To support the UK's AI and science ambitions, we must not only build more compute capacity, but also deepen national capability across software, skills, data and user support.

At the heart of this effort will be a new network of National Supercomputing Centres (NSCs). NSCs will anchor the UK's most powerful public systems while playing a far broader role as centres of excellence in the design, delivery, and effective use of compute. These hubs

will consolidate and uplift existing strengths, serving as focal points for collaboration between public bodies, universities, industry and research institutions.

NSCs will join up the compute ecosystem with the rest of the research infrastructure landscape – most importantly the UK's flagship programmes on dataset curation, such as the National Data Library (NDL), the new Health Data Research Service (HDRS), and the new OpenBind initiative.

Each NSC will be funded to support a diverse community of users from AI labs and scientists to SMEs and industrial partners. They will deliver high-performance systems, large-scale datasets, trusted research environments, and national programmes for software and skills. They will help ensure that compute infrastructure is not only accessible, but optimised and impactful, enabling users to maximise its potential.

While NSCs may differ in size and scope, they will share a common mission: to build a coherent, federated ecosystem where the UK's compute capabilities are connected, strategic, and globally competitive. They will also serve as innovation testbeds, trialling new architectures, incubating emerging technologies, and forming international partnerships to position the UK at the forefront of global scientific infrastructure.

The Edinburgh Parallel Computing Centre (EPCC) will be the UK's first National Supercomputing Centre. Home to the new national supercomputer and already a leader in High-Performance Computing (HPC) services, user support, and software development, EPCC is both nationally and globally recognised. EPCC already leads on the training of research technology professionals and, going forward, will play a leading role in the UK's next HPC ecosystem coordination consortium. We will designate additional centres to form a federated network of NSCs, collectively supporting the UK's long-term scientific, technological, and strategic ambitions.

Recognising that expertise will exist beyond the National Supercomputing Centres, we will establish a UK Compute Community Forum to bring together system operators, domain experts and end users. This group will help shape procurement priorities, surface emerging requirements, and ensure our infrastructure remains responsive to changing needs and technologies.

To support better system design, UKRI will also launch a Living Benchmarks Library – a dynamic, open-source suite of real-world workloads developed with input from researchers, vendors, and hardware innovators. This library will ensure future procurements are based on actual UK research and innovation needs, not theoretical metrics, and will help us stay open to new and disruptive technologies.

We will continue investing in research software and skills, launching a set of software challenges in 2026 and building on programmes for research software engineers (RSEs) and research technology professionals (RTPs). This will complement our broader programme of action to support AI training and skills in response to the AI Opportunities Action Plan.

Together, these efforts will ensure that the UK's compute ecosystem is effectively joined-up, supported by software, data, and skills – and remains responsive to the communities it serves.

Action 3 – Partner with like-minded countries and computing centres to expand access to a broader range of research infrastructure, facilitating collaboration and skills and knowledge exchange.

The most important breakthroughs – whether in climate science, AI safety, or life sciences – are rarely achieved alone. The best research is done in collaboration, and the UK will continue to be a committed partner in the global compute and research ecosystem.

Through our international partnerships, we aim to enhance access to compute capabilities and architectures not currently available in the UK, enable skills and knowledge exchange, and create new platforms for collaborative research. These relationships help ensure UK researchers, innovators and public institutions can remain at the cutting edge of AI and scientific discovery.

We also know that international collaboration is essential to tackling shared challenges, particularly in areas like security, resilience, and sustainability. Working with trusted allies will help develop common standards, jointly shape governance around AI compute, and drive collective innovation in areas such as clean energy integration, secure infrastructure, and responsible AI development.

We are already taking steps to maximise our membership of the EuroHPC Joint Undertaking, Europe's leading supercomputing partnership. The UK is committed to being a reliable and active partner, ensuring UK researchers, companies and public institutions benefit from shared infrastructure, reciprocal access, and joint missions.

UK researchers and businesses already benefit from access to some of the world's largest supercomputers, and new AI factories, through the UK's association with Horizon Europe and the EuroHPC membership. As we build out our sovereign domestic capabilities, we will take steps to further support researchers to fully participate in European projects and build vital connections and collaborations.

We are also exploring new bilateral and multilateral partnerships, such as our recent collaborations with Canada and France, and will seek to deepen ties with countries that share our values and ambition. To be most impactful, we will facilitate cooperation at every level – from government-to-government agreements to partnerships between research institutions, businesses, and start-ups.

In an increasingly volatile and competitive world, building deep, trusted international relationships will be essential not only for advancing science, but for securing our strategic position in the global AI landscape. We will lead where we can and collaborate where we should, ensuring the UK is embedded in the global networks that will shape the future of compute and AI.

2.3 The UK public research compute ecosystem of the future

Through these steps and across the AIRR and our National Supercomputing Centres, we will build out a research compute ecosystem that delivers on our priorities. It will be an ecosystem that meets researchers at their point of need and seamlessly guides users to the resources they need to progress their work.

We will achieve this by creating a dynamic research platform that provides users with the tools they need at all stages of the research lifecycle by integrating and federating a diverse range of compute solutions. We will adopt a singular and cross cutting approach to providing users with the skills, software and data they need – and enable them to access those resources at all stages of the research lifecycle. Finally, we will remain committed to co-designing and iterating our approach to ensure that as users' needs evolve, so too does the infrastructure that serves them.

3. Putting compute to work – powering innovation across the public and private sector

3.1 The case for change – moving to a more strategic and outcomes focused approach

The allocation of compute defines how this critical national asset is used. As compute becomes ever more central to science, AI, and economic growth, getting allocation right is more important than ever. It is no longer just a technical or operational question – it is a strategic one.

Historically, the UK's approach to public compute allocation has too often been slow and fragmented. Processes have prioritised system utilisation over outcomes, favouring high throughput rather than high impact. For many users, especially in AI, interdisciplinary science, and emerging industries – the system has been hard to navigate and unresponsive, limiting access and discouraging ambition. This must change.

Through a refreshed allocation model for the AI Research Resource (AIRR), we will create a system that is strategic, agile, and outcome-focused. Rather than simply filling machines, we will prioritise transformative research, national missions, and innovation with real-world impact. The process will be faster and more transparent, with clearer pathways for researchers and innovators – across both the public and private sector. It will be an allocation model that incentivises and rewards the most ambitious projects and those with the greatest potential to shape the UK's future.

We are also clear that this compute infrastructure will serve not just the UK's leading research community, but also be there for public sector organisations, start-ups, SMEs, and industrial users to innovate, experiment, test, and develop new products and solutions.

Action 4 – Introduce a refreshed allocation model that targets compute toward the UK's highest-impact research and innovation priorities, backing mission-driven research and strategic national goals.

The new allocation model will place greater emphasis on shaping science while continuing to ensure we serve the needs of our researchers. The following principles will guide our approach:

- **Backing our national priorities.** We will back the projects that matter most to the UK – those that align with national priorities and that deliver on our plan for change.
- **Big bets over small wins.** We will put our weight behind transformative, large-scale research programmes. In particular those that deliver real-world impact and deliver breakthroughs that change lives and grow the economy – cultivating a spirit of experimentation.
- **Moving fast, to ensure we don't miss out.** We will allocate compute with speed and precision. Where opportunities to back ambitious and transformative projects arise, we will act and break through bureaucracy.
- **Inclusivity, but not equality.** The AIRR will be open to all, but capacity will not be thinly spread. We will emphasise projects that have the most credible plans and clearest outcomes, with AI at their core.
- **Acting fast in times of crisis.** The AIRR is a national resource and we will use it to respond quickly to urgent or emerging needs, including public health threats and security scenarios.

The allocation model will not be static. It will adapt to the changing needs of users, advances in AI, and the expansion of compute capacity across the UK ecosystem. We will continuously refine our approach, ensuring that it remains responsive, impactful, and aligned with national priorities.

3.2 Allocation priorities

Capacity will be distributed across three broad access tracks:

- **General use, supporting a broad pool of public sector, academic and industry users.** This track will support shorter-term, project-based allocations, with a more accessible and agile application process to encourage new users, disruptive thinking, and early-stage innovation. The aim will be to provide a responsive and inclusive pathway for diverse research and development needs, enabling experimentation, collaboration, and capacity building across the ecosystem.
- **Mission-driven scientific research as set out in our upcoming AI for Science Strategy, such as AI for engineering biology, frontier physics including nuclear fusion, materials science, medical research and quantum technology.** Allocations will be made through structured calls, typically offering medium to long-term access depending on the scope and scale of the project. Each call will set out clear thematic priorities and areas of focus, aligned with national missions and emerging needs. We will encourage multidisciplinary and cross-sectoral proposals, recognising that the most ambitious and impactful projects often span traditional boundaries.

Action 5 – Guarantee dedicated compute access for Sovereign AI and the AI Security Institute, ensuring the UK’s core AI functions have the capacity they need to lead.

- **We will provide dedicated tracks for the governments core strategic AI functions – the AI Security Institute (AISI) and Sovereign AI Unit.** This will ensure that both functions have access to the sustained capacity they need to deliver on their missions.

The Sovereign AI Unit

Announced in the AI Opportunities Action Plan, the Sovereign AI Unit sits in DSIT – backed by up to £500 million in partnership with the British Business Bank – with a mission to take bold and rapid action to secure the UK’s future as a sovereign AI nation in the context of rapidly advancing AI capabilities. The Unit has started delivering with projects including:

- **OpenBind:** £8m seed investment in a Diamond Light Source (DLS)-led consortium to create the world’s largest open protein-ligand dataset and secure sovereign leadership in AI-native drug discovery
- **Partnerships:** with Anthropic and Cohere to gain a UK “stake” in frontier AI development, ensuring that strategically important companies are invested in the UK and that their presence is secure.
- **Encode Fellowship:** seeking to expand a flagship 1-year, entrepreneurially focused program run by Pillar VC and ARIA aiming to get top global AI talent working in UK labs on various domains in AI for Science within 4 months.

AISI will utilise capacity to support its critical AI safety and security work, including the evaluation of foundation models, red-teaming, and frontier risk analysis. AISI’s role in building national and international leadership in AI safety and security requires it to operate independently, flexibly, and at scale – supported by reliable compute infrastructure embedded within trusted environments.

For the Sovereign AI Unit, dedicated compute access for companies and projects it sponsors will be essential to supporting the UK’s ambition to lead in key AI priority areas – whether in emerging AI paradigms, high-impact sectors like healthcare or defence or through backing the most promising UK-based start-ups. Sovereign AI will use this capacity alongside its broader toolkit including access to high-value data assets to drive forward UK capability and accelerate AI development where it matters most. This approach reflects our clear commitment to using public compute strategically not just as infrastructure, but as a lever to deliver real outcomes for the UK economy and ensure British AI leadership in critical areas.

AISI and the Sovereign AI Unit will be responsible for allocating capacity according to its specific operational needs, mission objectives, and delivery plans. This approach ensures both flexibility and strategic alignment, empowering each organisation to focus on outcomes while maintaining accountability for how compute is used.

4. Building AI infrastructure to keep the UK at the cutting edge of AI development

4.1 Global and technological context

While public infrastructure plays a critical role in supporting national priorities and scientific leadership, it will make up only a small share of the UK's total compute capacity. The vast majority will come from private infrastructure, built to serve growing commercial demand for training, inference, and AI-powered services.

Here the landscape is evolving at speed. As AI models become more powerful and adoption ramps up, the demand for compute is growing exponentially. The compute needs of AI model training continue to scale, requiring specialised clusters made up of vast chip fleets with advanced networking and integrated energy systems. At the same time, inference is becoming increasingly critical, as AI is deployed across the economy and as new innovations move 'thinking' and therefore demand to the inference layer. Inference workloads are therefore expected to make up the majority of total demand and will be key to unlocking AI's full economic potential.

While the UK is home to a significant and mature data centre market, it is not yet optimised for AI – and certainly not at the scale required to support frontier training or national-scale inference. Most existing facilities are geared toward general-purpose enterprise computing, lacking the density, energy integration, and technical design needed for high-intensity AI workloads. Meanwhile, other countries are moving fast. The United States, United Arab Emirates, and several European nations are building out dedicated AI campuses that will each surpass 1GW of capacity.

4.2 Our approach – moving now to build future proofed capacity

Without urgent action, the UK risks being left behind – over-reliant on foreign infrastructure and missing the opportunity to embed AI capability into our economy and to anchor the next generation of AI companies here at home.

That is why now is the time to build. Infrastructure of this scale takes years to plan and deliver, requiring access to land and power, grid upgrades, and access to large amounts of capital. By moving decisively, the UK can establish itself as a leading location for AI infrastructure, attracting investment, supporting innovation, and building clusters of high-value activity. AI Growth Zones are our strategy for delivering this.

Action 6 – Deliver large scale AI Infrastructure via AI Growth Zones across the UK, ensuring the UK has the capability to support AI training and inference workloads.

Our approach to AI Growth Zones is focused on facilitating the delivery of the largest, most strategic infrastructure projects that would otherwise not be feasible in the UK. These are the trailblazing, high impact builds that will anchor the development and deployment of AI across the UK and attract global investment – they will be amongst the largest infrastructure projects of the next five years.

We forecast that the UK will need at least 6GW of AI-capable data centre capacity by 2030 – a threefold increase on the data centre capacity that is available in the UK today. Should the capabilities and adoption of AI accelerate, demand could exceed this baseline significantly.

Our goal is to establish a core group of nationally significant sites, each capable of serving at least 500MW of demand by 2030, with at least one AI Growth Zone scaling to more than 1GW by 2030. Ideally, each site will be designed to expand into the 2030s, so that we can add capacity as AI demand continues to grow.

AI Growth Zones will support a range of AI workloads, including both training and inference. We recognise that AI training is generally less dependent on location and that by contrast, inference may benefit from proximity to data sources or end users. Therefore, each Growth Zone will have its own architectural needs, and we will work with investors and developers to address these.

We aim to confirm and commence development of the first AI Growth Zones by the end of this year, delivering sites across the UK, including in Scotland and Wales. Each location will reflect local energy availability and regional industrial opportunities, and will be aligned to with our overarching growth, clean power, and decarbonisation objectives.

Noting that AI infrastructure is enormously energy intensive and demand is only increasing, powering the next generation of AI infrastructure sustainably will be one of the defining challenges of the coming decade. The government is committed to delivering AI Growth Zones sustainably, ensuring impact on energy, water and wider environmental systems are considered and minimised.

However, this challenge also represents an opportunity to pioneer new ways of integrating energy and infrastructure, and to lead the world in how to build sustainable, scalable systems from the ground up.

Action 7 – Explore new models for delivering the energy infrastructure that AI demands, including through renewables, advanced nuclear, and innovative grid solutions.

The UK has deep strengths in energy innovation. We are already a global leader in offshore wind, with a rapidly growing base of solar, battery, and flexible grid technologies. More recently, we have committed to a new generation of nuclear, including Small Modular Reactors (SMRs) that offer compact, clean baseload that is well suited for high-intensity workloads like AI compute. AI Growth Zones will provide the physical and policy environment to bring these technologies together in new, more integrated ways.

Therefore, to establish AI Growth Zones we will support the development of behind-the-meter, low-carbon energy solutions – designing campuses that are more self-reliant, easier to scale, and better aligned with local energy capacity. We will explore low-carbon onsite generation, microgrids, battery storage, and flexible demand systems.

In addition, we will explore options to deliver an AI Growth Zone that incorporates advanced nuclear capability, such as an SMR, demonstrating what a next-generation, low-carbon AI campus can look like.

Countries that can deliver this kind of infrastructure at the intersection of energy, AI, and advanced construction will not only meet their domestic needs – they will also lead in exporting the solutions, supply chains, and standards of the future.

4.3 Going beyond infrastructure to deliver widespread national and local benefits

Hosting large-scale AI infrastructure in the UK brings immediate and long-term national benefits. At the most basic level, building this infrastructure on UK soil strengthens the economy and enhances our strategic and economic resilience. AIGZs will ensure we are not reliant on overseas providers and insulate key elements of our economy from global shocks. In addition, by siting data centres in strategic locations, AIGZs will relieve pressure on congested grids, and support smarter regional infrastructure planning.

Our plans for AI Growth Zones go beyond infrastructure alone. They reflect the extraordinary interest across the UK - from industry, academia, and regional leaders - to participate in the development and deployment of AI. The UK brings together a powerful combination of world-leading research institutions, deep pools of talent, innovative businesses, and globally competitive sectors. AI Growth Zones offer a way to harness and coordinate these strengths - creating places where AI capability can thrive and where national ambition meets local opportunity. We are committed to ensuring AI Growth Zones deliver a wide range of national and local benefits.

Action 8 – Harness AIGZs to deliver both national and local benefits across the economy, serving as platforms for innovation, adoption and collaboration through which AI capabilities are developed, trialled, deployed and scaled.

We will pursue a multi-faceted approach to AI Growth Zones – in addition to strengthening the UK's security and resilience we will explore opportunities to deliver further AIGZs as platforms for broader innovation and strategic impact – including driving AI adoption across sectors, supporting cutting-edge research and development, and enabling experimentation with emerging AI paradigms such as embodied AI, autonomous systems, and human-AI collaboration. AIGZs will ensure all parts of the UK serve as engines, accelerating not just infrastructure build-out but the UK's capability to develop, deploy, and lead in next-generation AI.

These will include:

- **R&D and Innovation Platforms** – AIGZs that enable cutting and bleeding edge research pilots that push the boundaries of applied AI. This includes trialling advanced robotics, automated services, and new AI paradigms, as well as fostering collaboration between start-ups, research institutes, and industry. These zones will create the conditions to turn cutting-edge research into deployable solutions, cementing the UK's role as a leader in AI innovation.
- **Adoption Testbeds** – AIGZs that serve as real-world testbeds for AI adoption across sectors and regions. From healthcare and transport to energy and local government, these zones can pilot the integration of AI into critical systems and services - demonstrating impact, reducing barriers to deployment, and accelerating adoption. By working across local partners, these AIGZs will showcase how AI can directly improve lives, productivity, and public services.
- **AI Workforce and Talent** – to realise the full potential of AI, we must ensure people across the UK can develop, deploy, and work confidently with it. Therefore, as we deliver AI Growth Zones, we will work with industry, investors and local partners to deliver projects that support the upskilling of local workforces to adopt AI. Alongside this, we will also deliver wider initiatives as part of the AI Opportunities Action Plan to deliver training and capability-building programmes to support the development and deployment of AI.

5. Creating sovereign, secure and sustainable capability

In an era where artificial intelligence will underpin everything from economic competitiveness to national security, having sovereign capability in AI – and the compute infrastructure that powers it – is essential. Sovereign infrastructure allows us to protect sensitive data, direct resources toward national priorities, and ensure we remain resilient in the face of global instability and supply chain disruption. Beyond resilience, sovereignty is also about opportunity: it allows us to shape our own future, build competitive advantage, and grow high-value sectors that support jobs and innovation across the UK.

AI is one of the world's fastest-growing markets – and compute is its beating heart. The global AI chip market is projected to be worth hundreds of billions by 2030, while investment in AI data centres is expected to exceed \$1 trillion globally by the end of the decade. These are not just enablers of innovation, they are industries in their own right, with supply chains stretching from advanced semiconductor design and fabrication to infrastructure delivery and operations. Countries that move quickly to build sovereign capability across this stack stand to benefit economically and geopolitically.

The UK has real strengths to build on. We are home to ARM, one of the world's most successful semiconductor design companies, and a growing network of AI and chip start-ups pushing the boundaries of the possible – from new architectures to energy-efficient designs. Our research institutions, including centres of excellence in Cambridge, Bristol, Edinburgh, London, and beyond, are leading in everything from AI model development to quantum computing and photonics. By combining our world-class talent with strategic investment and a focused industrial policy, the UK can secure its place as a global leader – not just in using AI, but in designing and building the infrastructure that powers it.

That is why Advanced Connectivity Technologies, Cyber Security, Semiconductors, and Quantum Technologies are priority growth markets in the UK's Modern Industrial Strategy. With our investments in AIRR and AIGZs, we recognise that compute infrastructure is the key market to drive sovereign capabilities and economic growth in these technology sectors.

5.1 Our approach to sovereignty – pragmatic, balanced and strategic

The UK's approach to sovereignty in compute is pragmatic, strategic, and uniquely shaped by our place in the world. As a deeply connected country with long-standing partnerships with the United States, the European Union, and other global allies, we do not define sovereignty as isolation or self-sufficiency.

Instead, for the UK, sovereignty means ensuring we have the ability to act independently and effectively where it matters most – to allocate compute to national priorities, protect sensitive data, and support UK research, innovation, and public services on our own terms.

In practice, this means ensuring that programmes like the AI Research Resource and AI Growth Zones result in capacity that is physically located in the UK and strategically aligned with our national goals – not just capacity rented from abroad.

We will continue to work with the world's best suppliers and collaborators as we deliver these programmes, maintaining access to trusted, secure, UK-based infrastructure that we can scale and govern as needed. But as the UK builds out its national compute infrastructure through a renewed public research ecosystem and the delivery of AI Growth Zones, we have a unique opportunity to use this platform to pull through homegrown technologies and solutions – providing a launching pad for British excellence.

5.2 From prototype to deployment – a pull through mechanism for UK suppliers

Action 9 – Support British companies to develop sustainable and secure compute technologies, using AIRR and AI Growth Zones to provide access, testbeds, and opportunities to scale.

The UK has both a rich heritage in chip design and a growing ecosystem of early-stage innovators ready to compete in the computing paradigms of the future. The next generation of compute infrastructure will need to be more secure, more efficient, and more sustainable. This will entail specialised chips, quantum systems, neuromorphic architectures, and memory-intensive designs – that go beyond traditional CPU/GPU configurations that dominate the market today.

Our goal is to create a pull-through mechanism that links the UK's world-class research base with testbeds and real-world use cases, enabling the most promising British technologies to scale and succeed. We will start by leveraging the UK's growing ecosystem of early-stage support, including ARIA's Scaling Compute programme, ChipStart UK, the new National Semiconductor Centre, and the next phase of the National Quantum Technologies Programme, and build on this with further strategic programmes supporting novel compute technologies. These initiatives are backing the development of low-energy, high-performance, and novel architectures – but they need access to national testbeds to prove performance and integrate with real-world systems.

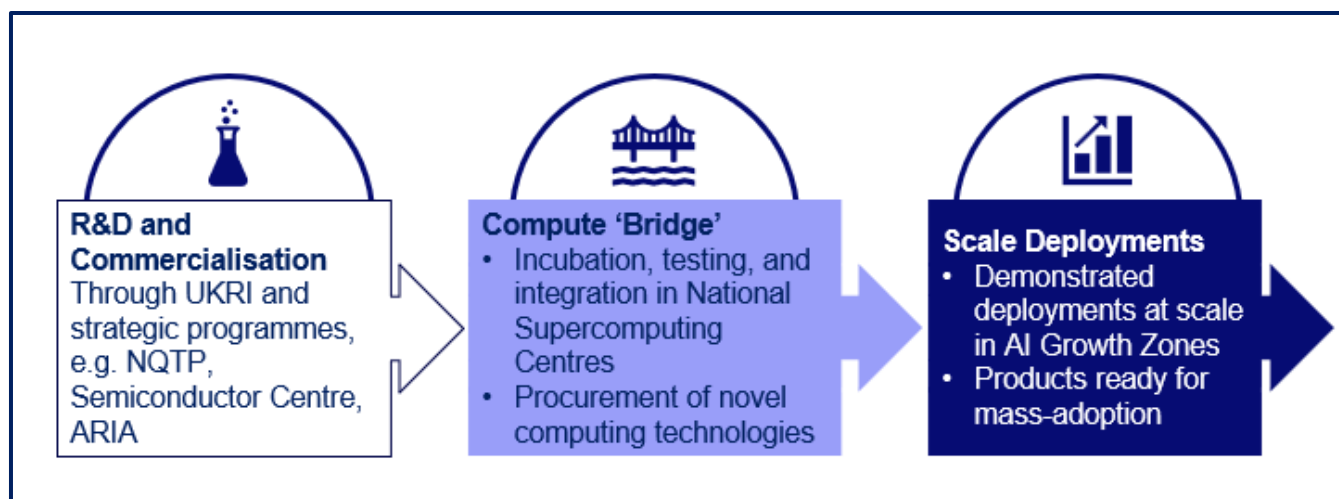
We will leverage our investments in the public compute ecosystem and purpose-built environments of National Supercomputing Centres to act as a trusted proving ground to test, validate, benchmark, and deploy emerging technologies. In particular, we will support

energy-efficient solutions that deliver greater compute efficiency for key AI workloads and highly secure by design architectures that strengthen cyber security and resilience.

In addition, through UKRI's living benchmarking exercise, we will improve the eligibility of novel compute technologies for procurements. This will ensure that innovative and emerging technologies are not discounted from procurement exercises because they do not meet criteria set to evaluate performance against incumbent solutions.

Our ambition is to pull through the most successful UK-developed technologies into the commercial, at-scale deployments of AI Growth Zones. These zones will offer the scale, power, and integration needed to showcase British capability on a global stage. In an ideal end state, we envision an AI Growth Zone powered by a full UK-designed compute stack – from chip to system to software – demonstrating what sovereign capability truly looks like and opening up new export and growth opportunities for UK firms.

Figure 2: The Compute 'Bridge' as a pull-through mechanism for UK compute tech



5.3 Compute – An opportunity area for Sovereign AI

The Sovereign AI Unit was announced in the AI Opportunities Action Plan as a new initiative specifically designed to build homegrown capability in emerging areas of the AI ecosystem. Backed by £500 million at the Spending Review, the programme is focused on unlocking opportunities where targeted public intervention can support UK leadership, scale national champions, and secure long-term strategic advantage.

Sovereign AI will focus its support on projects that meet a clear set of strategic criteria. In particular, it will prioritise initiatives that drive novel, IP-rich capability advancements, align with emerging technical and market trends in AI, and serve as a source of strategic influence.

Action 10 – Establish compute as a priority area for the Sovereign AI Unit, supporting research into new computing paradigms and helping British companies grow into global leaders

Compute as a priority area for Sovereign AI. This reflects the UK's unique strengths in chip design, software optimisation, and emerging architectures, as well as the strategic importance of securing access to the infrastructure that underpins AI development. Sovereign AI will build on its reserved allocation within the AIRR, using this capacity as a platform to support innovation across the compute stack.


Sovereign AI will retain the flexibility to move quickly, offering strategic backing, capital, and infrastructure access to accelerate promising technologies. In doing so, we will help ensure that British companies are not just users of global compute systems, but builders of the next generation of sovereign, secure, and scalable infrastructure.

While building homegrown capability will take time, this roadmap sets out our clear intention to invest in and grow sovereign strength across the UK's compute ecosystem. We recognise that this is a long-term effort – but one that starts now. We look forward to working closely with the UK's vibrant network of researchers, start-ups, industry leaders, and infrastructure providers to realise this ambition and ensure that British innovation powers the future of AI.

6. Conclusion and next steps

The launch of the AI Research Resource is a first step in delivering this Roadmap. We will now build on this momentum to realise our long-term vision, seeking to make quick progress as we deliver against our comprehensive plan. Our approach will be collaborative and inclusive: we will work with researchers and start-ups, local government and devolved administrations, industry and international partners to build a truly national and global ecosystem. And as AI and compute technologies continue to evolve at pace, we will remain agile and responsive, ensuring our strategy stays current and dynamic – always focused on driving innovation, increasing productivity, and unlocking sustainable growth for the UK.

7. Compute Roadmap – at a glance





Deepen the Ecosystem

- Recognising further NSCs and clusters of expertise in supercomputing
- Strengthen platform-level support for software engineering, data stewardship, and technical skills
- Expand international partnerships



Integrate the Future

- Begin integration of novel computing technologies – including **quantum, neuromorphic, and AI-hybrid** systems – into public infrastructure

2030+

Refresh, Reinvest, Reinvent

Renew commitments in the next Spending Review



Refresh the Roadmap

- Full strategic refresh of the roadmap in 2030 evaluating delivery, impact, and emerging global trends



Continue to Build

- Further scaling and development in AI Growth Zones, evolving capabilities to meet emerging needs



Replace and Review

- Begin procurement of a successor to the Edinburgh supercomputer with **no loss of service**.



Adapt to a Changing World

- Reassess user needs across science, public sector, and industry
- Integrate new computing paradigms into investment roadmap, respond to future missions, and maintain the UK's strategic edge.

E03400533

ISBN 978-1-5286-5894-2

This publication is available from: www.gov.uk/official-documents.

If you need a version of this document in a more accessible format, please email alt.formats@dsit.gov.uk. Please tell us what format you need. It will help us if you say what assistive technology you use.