COMMONWEALTH SECURITY & RESILIENCE GROUP

FIOR

The Enforcement Layer for AI

COMMONWEALTH AI SECURITY

POLICY AND DEPLOYMENT

Sovereign AI across the Commonwealth

FIOR Group • Commonwealth Security & Resilience Group

Joint discussion paper | November 2026

Prepared in connection with the Commonwealth AI Security Summit, London, 24–25 November 2026.



Foreword by Commonwealth Security & Resilience Group

As artificial intelligence becomes integral to the functioning of our economies and institutions, how much control do we retain over the technologies on which we increasingly depend? Against a backdrop of growing geopolitical uncertainty, governments across the Commonwealth are considering how they might securely diversify their supply chains while at the same time developing their own sovereign capabilities. The concentration of advanced AI development in the hands of a few companies presents a formidable challenge to national sovereignty, given the scale of both the economic opportunities and the security risks associated with the technology.



"Cyber security is a shared responsibility, and strong partnerships are essential to defending our networks in an increasingly contested digital environment."

Major-General Dave Yarker, Canadian Armed Forces Cyber Commander

What, then, are the risks for countries that fail to secure greater sovereign access to, and control over, AI? Dependence on overseas suppliers represents a significant source of political and economic vulnerability. Countries risk becoming susceptible to external pressure and coercion in areas over which they have historically exercised autonomy, while access to capabilities essential to national industries could be restricted or withdrawn.

For advanced economies, the stakes include the security of critical infrastructure, industrial competitiveness, and the protection of intellectual property. Emerging economies risk becoming major consumers of AI without deriving a commensurate share of its economic benefits. Across much of the Commonwealth, and particularly among its developing economies, the provision of essential public services could become increasingly dependent on commercial and political decisions made well beyond national borders. Nor can these risks be considered in isolation. Given the interconnected nature of global supply chains, the Commonwealth's more technologically advanced economies have a direct interest in the security and resilience of AI systems in use across its emerging and developing economies.

But what does it mean, in practical terms, to have sovereignty over AI? Even the world's largest economies rely on international supply chains for critical elements of AI development and deployment.

© 2026 FIOR Group Limited
Visit our website: fior.group

Sovereignty need not imply technological self-sufficiency.

What matters is the degree of control a country or organisation retains over the technologies it utilises: its ability to decide which technologies it acquires and how they are used, to protect sensitive information and to enforce its laws.

Even an AI service procured from a domestic or trusted supplier may rely on numerous overseas models, software providers and data processors. Without sufficient knowledge of these relationships, governments and businesses may inadvertently relinquish some of the very sovereignty they are seeking to establish.

The question of sovereignty assumes particular significance when we consider the security implications of AI. The protection of citizens, territory and national institutions has long been understood to be among the principal duties of the sovereign state. As we entrust increasingly autonomous AI systems with sensitive information and greater decision-making authority, we ought to reckon with the possibility of losing control over their actions. How can Commonwealth governments exercise meaningful oversight over technologies developed outside their jurisdiction? And how can they ensure that the authority they delegate to AI remains theirs to withdraw?

"Our defence force is integrating AI into our daily operations, but the level of integration is contingent on the existence of digital architecture to support the use of digital technology. With that in mind, the process of digitising the TTDF is ongoing and, for that reason, building relationships with trusted supply chains and partners is critical for operational security, effective service delivery, and business continuity."

Commodore Don Polo, Chief of Defence Staff, Trinidad and Tobago Defence Force



Few Commonwealth countries can hope to address the full range of these challenges alone, but there is much we can do in concert. The development of sovereign capabilities in AI security is a more attainable ambition for most Commonwealth countries than the development of frontier AI models. Longstanding relationships of trust across the Commonwealth lend themselves to a reciprocal exchange of expertise that might inform national policymaking. In law, our largely shared common law heritage provides a useful basis for cooperation as legislators and courts confront new questions of liability and accountability. In technology, we can draw on the considerable industrial diversity of the Commonwealth and some of the world's leading centres of innovation.

How, then, can the Commonwealth draw on its collective capabilities to help uphold the sovereignty of all members? The following paper examines what meaningful

© 2026 FIOR Group Limited
Visit our website: fior.group

sovereignty over AI entails in practice and how Commonwealth countries might achieve it, both individually and cooperatively, within their technological and economic means.

Anthony Halley | Director, Global Strategy and Programmes, Commonwealth Security & Resilience Group

Introduction | David Williams, Founder, FIOR Group



The Commonwealth brings together countries with profoundly different economic resources, technological capabilities and public-service needs, yet a shared interest in making artificial intelligence useful, secure and accountable. FIOR's conviction is that practical cooperation on AI security and sovereignty should begin with the institutions that will actually deploy these systems.

AI is becoming an operational actor, not merely a source of information. It will increasingly recommend decisions, access records, use software tools and carry out delegated tasks. That creates opportunities for growth and better public services, but also a fundamental question of authority: who decides what an AI system may do, how is that decision enforced, and can a nation retain control when external connectivity or commercial arrangements fail?

Sovereignty does not require each country to invent every chip, model and application. It requires a deliberate choice about critical dependencies, custody of data and intellectual property, continuity of essential services and the enforceability of human authority. There is substantial room for Commonwealth countries to cooperate on designs, procurement, independent security assurance and technical skills while preserving their separate national choices.

This paper proposes a practical path: deploy models that are economical and suitable for the task; keep sensitive adaptation and operation within the required jurisdiction; map delegated authority before agents are released; enforce that authority at runtime; and assemble evidence that the system will remain secure and useful. The UK can contribute particularly through its research, emerging hardware industry and security-assurance institutions, while fellow Commonwealth members should be equal partners in defining standards and creating locally owned capabilities.

David Williams | Founder, FIOR Group

© 2026 FIOR Group Limited
Visit our website: fior.group

1 | Why AI sovereignty and AI security now matter

Sovereign AI is the ability of a state or institution to determine how an AI capability is built, governed, operated, changed and, if necessary, disconnected from external suppliers. This is a question of effective control, rather than the nationality of every component.

Security is essential to the definition: custody of weights or data alone does not prevent an autonomous agent from undertaking an action outside its delegated authority.

AI systems are becoming embedded in education, public administration, healthcare, financial services, national infrastructure, defence support and private enterprise. Access to a foreign-hosted model can accelerate adoption, but its operation may depend on decisions taken by an overseas provider: pricing, model replacement, suspension, remote administration, regulatory access, retention policies, and the availability of data-centre capacity. Some dependencies are acceptable; for critical services, they must be visible and manageable.

AI is becoming an operational actor, not merely a source of information. It will increasingly recommend decisions, access records, use software tools and carry out delegated tasks. That creates opportunities for growth and better public services, but also a fundamental question of authority: who decides what an AI system may do, how is that decision enforced, and can a nation retain control when external connectivity or commercial arrangements fail?

The security challenge is changing as well. A conventional application normally executes an authorised sequence of commands. An agentic system may plan its own steps, call external tools, combine information across institutional boundaries, and retry after failure. A valid model response or an authenticated agent identity does not establish authority to act. Security therefore needs to cover the intended outcome, the identity of the agent, the provenance of the instruction, the scope and duration of delegated authority, the tool call and the resulting audit evidence.

The goal is assured national choice: the ability to move models, change suppliers, retain sensitive data, operate offline when required and stop unauthorised actions before they occur. A Commonwealth framework should offer routes to these outcomes that are affordable for both large countries and small states.

© 2026 FIOR Group Limited
Visit our website: fior.group

2 | Background: AI across the Commonwealth

The Commonwealth comprises 56 independent member states with approximately 2.7 billion people, more than 60% of whom are 30 or under. Thirty-three members are classified as small states. These facts suggest both substantial long-run demand for AI-enabled services and wide variation in domestic compute capacity, fiscal room and technical skills. They should not be mistaken for a measured Commonwealth-wide AI adoption rate. [CW1][CW2]

A single industrial prescription would be inappropriate. Large economies may develop substantial elements of the technology stack, while smaller states may achieve meaningful sovereignty through trustworthy local operations, procurement rights, regional infrastructure and shared evaluation capacity. [CW1]

AI activity spans different stages of maturity. The UK, Canada, India, Singapore and Australia have substantial research, enterprise or public-policy ecosystems; countries in Africa, the Caribbean and the Pacific are developing local innovation, language technologies and public-sector use cases under materially different infrastructure and skills constraints. Adoption statistics from individual national surveys are not directly comparable, and there is no robust single percentage of Commonwealth organisations using AI. A credible baseline should therefore distinguish access, pilots, production use, autonomous tool use and high-consequence use.

The Commonwealth Secretariat's 2026 resilience analysis identifies AI and disinformation as emerging pressures on public trust and democratic safeguards. For small states, concentration of infrastructure, overseas hosting, limited bargaining power and disaster resilience can make the ability to switch or operate locally more valuable than national ownership of a frontier-model laboratory. [CW3]

Commonwealth context	Implication for sovereign AI
Large technology economies	Invest selectively across research, compute, chips, models and shared assurance
Mid-sized and emerging economies	Focus on sector-specific models, skilled operators, local adaptation and procurement leverage
Small states and island economies	Prioritise resilient services, regional collaboration, open interfaces and tested supplier exit
All members	Adopt evidence-based security, human authority and cross-border interoperability

© 2026 FIOR Group Limited
Visit our website: fior.group

3 | A practical definition of sovereignty

It requires a deliberate choice about critical dependencies, custody of data and intellectual property, continuity of essential services and the enforceability of human authority. There is substantial room for Commonwealth countries to cooperate on designs, procurement, independent security assurance and technical skills while preserving their separate national choices.

A procurement specification should distinguish four questions: where data and derived artefacts reside; who controls operations, keys and administrative access; what rights exist to change and retain models and software; and whether the service continues during supplier failure or disconnection. Legal data residency, operational sovereignty and domestic ownership of intellectual property are different attributes.

A service can have some without all.

Dimension	Decision the UK customer must control	Evidence to request
Data	Where all sensitive inputs and derived artefacts are processed	Data flow map covering logs backups caches and support
Model	Use modification retention and replacement of weights	Licence rights provenance and exportable artefacts
Operations	Administration identity keys policy and release approval	UK operator model access records and key custody
Continuity	Operation during disconnection supplier failure or suspension	Offline test recovery exercise and exit plan
Security	Which agents may perform which actions	Enforced permissions denial tests and local audit
Supply chain	Which overseas dependencies remain acceptable	Component register contractual protections and alternatives

For a member state, replace each UK-only geographic condition in this matrix with the approved national or regional jurisdiction. Where a regional or allied cloud is acceptable, the procurement should make that delegation explicit. No claim of “sovereignty” should conceal operational control exercised through an external management plane.

© 2026 FIOR Group Limited
Visit our website: fior.group

4 | Middle powers, emerging economies, developing markets and Commonwealth cooperation

The argument that middle powers can cooperate to retain agency in a technology market dominated by a small number of suppliers is persuasive, provided cooperation is practical rather than symbolic. Shared reference architectures, standards, assurance evidence, model evaluation methods, pooled demand and procurement clauses can give participating countries more credible choices without requiring political or technological uniformity.



"There is a genuine opportunity for mutual advancement in diplomacy and trade through collaboration amongst Commonwealth Nations in Secure Sovereign AI."

Gemma Ungoed-Thomas. Former Director Cyber Security, Cabinet Office

The Commonwealth is neither a single market nor a security alliance. Members differ in law, national interests, procurement practice and threat assessment. A successful initiative should be voluntary, modular and based on reciprocal participation:

common technical tests and recognised evidence, with each state deciding its domestic control boundary and permissible partnerships.

The case is the same for both emerging economies and developing markets in the Commonwealth. While some regional groupings pair Commonwealth countries together in close economic and security zones, including in Africa and the Caribbean there is no singular agreed approach, and expertise, capital and resources might be found outside of these groupings. This is what makes Commonwealth cooperation particularly interesting. Within the wider network there exists the full value spectrum.

© 2026 FIOR Group Limited
Visit our website: fior.group

Illustrative Partner Grouping	Potential contribution to a Commonwealth approach
AI research, security evaluation and assurance (e.g. UK, Canada, Australia)	Advanced AI research and model science; AI and chip design; security evaluation and assurance institutions; sovereign reference architectures; enforceable agent security; model compression and optimisation; trusted deployment and advanced security testing.
Engineering, integration and deployment at scale (e.g. India, Malaysia, Bangladesh)	Large-scale software engineering and systems integration; cost-effective AI development and delivery; sector-specific deployments; semiconductor and electronics manufacturing; language and data diversity; adaptation and scaling of AI systems across different operational environments.
Financial, regulatory and digital governance expertise (e.g. Singapore, UK, Mauritius)	Financial-sector AI risk management; applied AI governance; rigorous regulatory and operational implementation; digital financial services; investment capital; international commercial frameworks; trusted cross-border AI services.
Critical infrastructure and operational security (e.g. Australia, UK, Singapore, South Africa)	Defence-oriented and critical-infrastructure AI deployment; cybersecurity and operational security guidance; secure integration with existing systems; protection of energy, transport, telecommunications and financial infrastructure; resilience testing and incident response.
Emerging technology markets and large-scale adoption (e.g. Nigeria, Kenya, India, Bangladesh, Ghana, Rwanda)	Large and growing populations and consumer markets; rapidly expanding digital economies; opportunities for cost-effective deployment and commercial scaling; entrepreneurial technology ecosystems; expanding technical workforces; development of AI applications across public services, finance, agriculture and industry.
Critical resources, energy and industrial capacity (e.g. Zambia, South Africa, Mozambique, Tanzania, Australia, Canada)	Critical minerals and industrial inputs supporting semiconductor, electronics and energy supply chains; opportunities for refining, processing and manufacturing; renewable and conventional energy resources; potential for investment in data centres, computing capacity and supporting digital infrastructure.
Strategic connectivity and regional infrastructure (e.g. Singapore, Malaysia, Kenya, Fiji, Malta, Jamaica, Barbados)	Strategic locations along major maritime and communications routes; subsea cable connectivity and diversification; secure regional data infrastructure; logistics and digital trade corridors; opportunities for distributed computing, regional digital services and more resilient supply chains.
Diverse operating environments, regional resilience and sovereign deployment (e.g. Rwanda, Kenya, Ghana, Barbados, Jamaica, Fiji, Samoa, Papua New Guinea)	Local priorities, languages and data governance; culturally and linguistically diverse environments for AI development and validation; testing under constrained connectivity and computing resources; disaster resilience and continuity planning; practical applications in maritime security, fisheries, climate resilience and public administration; opportunities for local ownership, trusted infrastructure and sovereign control of sensitive data.

© 2026 FIOR Group Limited
Visit our website: fior.group

5 | The UK's contribution: hardware, design principles and assurance

A credible opportunity in more of the hardware stack

The UK has become more ambitious about the portions of the AI technology stack it can credibly produce. It has strengths in processor intellectual property, chip design, systems engineering, research and emerging specialised inference hardware. Low-cost, high-throughput token generation is a particularly promising area, because inference economics increasingly depend on moving model weights and intermediate data efficiently, rather than on general-purpose compute performance alone. Success would allow UK-designed technology to serve selected sovereign workloads, including compact models and low-latency inference, and to capture more domestic economic value.

The UK AI Hardware Plan, published in June 2026, explicitly supports domestic chip capability. On 11 November 2026, the UK government and British Business Bank announced investment in Fractile, whose architecture seeks greater efficiency and throughput for AI inference. These are credible industrial developments, though pre-commercial performance claims must be validated in production under equivalent workloads. [CW4][CW5]

NVIDIA nevertheless retains decisive advantages in installed capacity, accelerator performance across diverse workloads, software tooling, developer adoption, manufacturing relationships and system-level integration. Recreating that complete ecosystem would be exceptionally costly, and UK-designed chips will still generally rely on international fabrication, memory, packaging or other supply chains. The rational strategy is therefore to build competitive UK capabilities in segments where architecture provides an advantage, while retaining access to NVIDIA and other global suppliers. Sovereignty can be strengthened by multi-vendor support and credible switching rights, not only by domestic substitution.

Sharing architectures and independent tests

The UK can help other Commonwealth members by sharing open, adaptable design principles for custody of data and weights, sovereign adaptation pipelines, key ownership, local policy enforcement, offline operation, interoperability and supplier

© 2026 FIOR Group Limited
Visit our website: fior.group

exit. These principles should be technology-neutral, with representative implementation patterns for cloud, premises and portable deployment.

The UK AI Security Institute (AISI) undertakes research and evaluations of advanced AI risks. The National Cyber Security Centre (NCSC) provides distinct cyber-security guidance and assurance arrangements, including the Cyber Resilience Testing (CRT) process under Principles Based Assurance. Model evaluations and product cyber-resilience assessment answer different questions; no programme should be described as certifying a complete deployment, or as automatically transferring a UK assessment into another country's accreditation. [CW6][CW7]

A Commonwealth programme could share test methodologies, train independent assessors, publish common evidence templates and enable reciprocal understanding of assessment results, while national authorities retain final certification and procurement decisions.

© 2026 FIOR Group Limited
Visit our website: fior.group

6 | Governance mapping: from institutional authority to machine-enforceable policy

Governance mapping should become a core competence in successful AI adoption, and a distinct professional service rather than an afterthought to installing an AI gateway. Complex organisations rarely possess one authoritative map of what every employee, automated process and external counterparty is permitted to do. Agentic systems make those gaps visible: authority may be delegated by a board, statute, department, contract, role or individual, and may be limited by monetary value, data type, geography, time, purpose or required review.

A governance-mapping engagement begins with business process engineering. It identifies the task, the accountable person, the information and systems touched, the decision rights involved, downstream consequences, other agents and organisations in the workflow, and the circumstances requiring escalation. It then reconciles existing policy with actual operating practice, records unresolved conflicts and specifies an enforceable decision at each consequential tool or API boundary.

Stage	Deliverable	Security outcome
1. Process discovery	Task and decision inventory; data and tool flow diagrams	Visibility into the full agent action surface
2. Authority analysis	Human principals, delegated mandates, constraints and counterparties	Clear source and limits of authority
3. Risk classification	Sensitive data, critical actions, segregation of duties, escalation thresholds	Proportionate controls
4. Policy engineering	Machine-readable allow/deny rules, approval paths and exceptions	Pre-action enforcement
5. Validation	Scenario tests, bypass tests, audit and revocation evidence	Demonstrable coverage and accountability
6. Continuous governance	Version control, recertification, incident learning and ownership	Policy stays aligned with institutional change

© 2026 FIOR Group Limited
Visit our website: fior.group



"Where we interoperate with allies and partners, there must be a common method for mapping and enforcing agentic processes."

Lt. Gen. Sir Tom Copinge-Symes Former Deputy Commander, Strategic Command

This requires consultants with experience implementing AI across sectors, strong knowledge of models and techniques, business-process redesign, senior-client relationship management and the ability to build a repeatable international advisory practice. FIOR's governance advisory capability will translate those maps into policy configurations implemented by the FIOR AI Gateway, while keeping design and assurance work separable from the technology purchase.

A mapping exercise must include machine-to-machine and organisation-to-organisation relationships. An agent working for a bank may interact with an external settlement provider, regulator or outsourced processor; each boundary requires clarity about delegated rights, credential ownership, data minimisation and evidential responsibility. An agent should not inherit broad human permissions simply because it acts on a human's behalf.

Mapping is not a one-off diagram. It should generate a maintained authority register, policy decision records, test cases, accountable owners and a controlled change process. Human oversight then becomes a concrete, auditable part of system design: approvals can be required for specific categories of action, and revocation can be enforced locally at machine speed, subject to the timing limits of disconnected operation.

© 2026 FIOR Group Limited
Visit our website: fior.group

7 | Security assurance in practice: FIOR's assessed gateway

Cyber security assurance and model assurance serve different purposes

We recommend assurance of both the software that operates an AI system and the model that produces its outputs. Cyber security assurance examines the engineering and controls that protect the product and its environment. Model assurance examines whether a particular model is suitable for its intended role, including accuracy, robustness, data provenance, harmful behaviour and performance under representative conditions. Neither assessment substitutes for the other.

This distinction is especially important for compressed and distilled models. Evaluation should establish which capabilities have been retained, where performance has changed and what operating limits apply. We recommend versioned evidence linking the base model, adaptation data, delivered weights, evaluation results and approved use cases. Material changes to models, tools or operating conditions should trigger proportionate reassessment. The UK AI institute's published evaluation approach also recognises that evaluations measure specified properties and have limitations. [21]

The UK can differentiate internationally by combining efficient models with assured software and enforceable authority over agent actions. This gives buyers a practical basis for trusting a sovereign service: evidence about the model's suitability, evidence about the surrounding product's cyber resilience, and controls governing what the resulting system is permitted to do.

The CRT assessment process and our completed assessment

Cyber Resilience Testing, or CRT, uses the NCSC's Principles Based Assurance methodology and its Cyber Resilience Testing Assurance Principles and Claims. A vendor works with an NCSC-assured Cyber Resilience Test Facility, providing engineering documentation and evidence of testing and operational processes. The facility evaluates the evidence against the relevant claims, identifies gaps and associated risks, and produces an assessment report in accordance with NCSC guidelines. CRT is a risk-based assessment rather than a simple pass or fail certification.

FIOR has become NCSC Cyber Resilience assessed. The resulting report confirms that all 72 assessed software security claims have been met.

© 2026 FIOR Group Limited
Visit our website: fior.group

This provides a strong foundation for our AI Gateway and for customers considering our enforcement technology within sovereign AI deployments.

The assessment methodology is time and cost effective and should serve as a precedent form Commonwealth partners.

© 2026 FIOR Group Limited
Visit our website: fior.group

8 | Sovereign deployment economics and delivery architecture

Frontier models: where is national investment justified ?

The economics of frontier competition differ from the economics of using AI. A training programme buys an uncertain chance of capability leadership at a particular point in time. A deployment programme buys an operating service. Confusing them can consume the budget intended to make hospitals, government departments and infrastructure operators more productive.

Epoch AI's 2024 research estimated that the amortised hardware and energy cost of final frontier training runs had grown about 2.4 times annually since 2016. Its projection of billion-dollar runs by 2027 was conditional on that trend continuing, rather than an audited statement of present spending. The distinction matters, but the research still illustrates the scale and trajectory of the commitment. [6]

Meta's account of Llama 3.1 describes training the 405 billion parameter model on more than 15 trillion tokens using over 16,000 H100 GPUs. This is evidence of the infrastructure involved in one major programme, not a minimum specification for every useful model. [7] The UK would also need to fund failed experiments, post-training, evaluation, data preparation and long-term maintenance. Hardware purchasing creates further obligations for power, cooling, networking, utilisation and replacement.

Benefits do not automatically accrue to the state

A domestic frontier model could create research spillovers, skilled employment, bargaining leverage and a strategic fallback. Those benefits deserve assessment. However, an expensive model does not automatically become competitive commercially, capture a user base or remove dependencies on imported accelerators and software. Funding must be evaluated against a realistic distribution strategy and the opportunity cost of the same resources elsewhere.

For many government and enterprise tasks, the operational bottleneck is access to accurate records, integration with legacy systems, workflow redesign or permission to take action. A more capable general model cannot resolve weak data stewardship or missing authority. Investment in a reusable national data pipeline or secure execution layer may support many models and institutions, while investment in one model family can concentrate risk.

The comparison also changes as external capability becomes cheaper. Stanford's 2025 AI Index reports a greater than 280-fold decline between November 2022 and

© 2026 FIOR Group Limited
Visit our website: fior.group

October 2024 in the price of inference at a specified GPT-3.5-level benchmark. This historical benchmark does not establish equivalent quality across tasks or guarantee future price falls. It does show why a public investment case should stress-test rapidly improving alternatives. [8]

The recommended investment boundary

A member country should not make matching the largest general purpose frontier LLMs its principal sovereign deployment strategy. The expected incremental benefits for routine sovereign workloads are unlikely to justify the full programme cost where smaller, adaptable models can meet the specification. Public buyers should require an explicit demonstration of unique strategic benefit before funding from-scratch general frontier competition.

There are credible exceptions. Fundamental research, AI security evaluation, scientific discovery, national language and public-interest datasets, and a contingency capability for high-consequence tasks can justify targeted support. Research compute has uses beyond frontier LLM training. The policy should preserve the ability to evaluate and develop models, while avoiding an open-ended obligation to remain on every global leaderboard.

An appropriate investment test asks whether a proposed model delivers a capability unavailable through licensed weights, distillation, retrieval or trusted partnerships;

whether that capability changes a material national outcome; and whether the member state can sustain the programme through several generations. If those conditions are absent, efficient adaptation and deployment should normally receive priority.

Compact, distilled and compressed models

Different techniques solve different problems

Knowledge distillation transfers aspects of a teacher model's behaviour into a student model using examples, predictions or other training signals. It can produce a model focused on a narrower task at a lower operating cost. Distillation remains a training activity: generating teacher outputs, preparing data and validating the student consume resources.

Compression is broader. It can remove structures or redundant parameters; quantisation reduces numerical precision. Fine-tuning adapts behaviour, while retrieval augmented generation supplies selected external information at inference

© 2026 FIOR Group Limited
Visit our website: fior.group

time. These techniques can be combined, but they are not interchangeable. Retrieval can keep a knowledge base current without repeatedly embedding every update in the weights, although access control must also apply to retrieval.

FIOR partnered with Multiverse (EU HQ but with an established Canadian office), whose published Compact if AI methodology describes profiling, structured pruning, healing through distillation and post-training that can include quantisation. It reports reductions of up to 95% for selected applications. The achievable reduction depends on the selected model and application. We assess the resulting service against its workload requirements; a reduction in model size does not translate directly into the same percentage reduction in total cost, energy or latency. [9]

Where the advantages arise

Smaller models can reduce the memory needed to hold weights and may lower accelerator requirements. This can change the deployment class from a large GPU cluster to a single server, workstation or portable device. It can also make replication across sites economical and reduce the distance sensitive information must travel. These are mechanisms for savings, rather than guaranteed outcomes.

A simple memory illustration makes the opportunity concrete. Eight billion parameters stored at four bits require approximately 4 GB for raw weights; 70 billion require about 35 GB. Runtime buffers, quantisation metadata, the key-value cache and the operating environment require additional memory.

Long contexts and concurrent users can dominate the remaining budget. Parameter count alone therefore cannot establish that an application fits a particular device.

Compact models are strongest where a buyer can define a clear acceptance boundary: extracting fields from standard documents, classifying incidents, retrieving approved procedures, drafting bounded reports or selecting from a limited tool set. They can be combined with deterministic rules and human review. Broad reasoning, unfamiliar coding problems, complex multimodal interpretation or long-context synthesis may still justify a larger model.

Distillation can inherit errors and unwanted behaviour from the teacher. Compression can degrade rare but important capabilities even when average benchmark scores remain strong. The evaluation set must cover sensitive subgroups, unusual cases, refusals, prompt injection, tool use and the actual language of the service. A small model is not intrinsically safer; its reduced footprint is valuable only if the resulting service remains fit for purpose.

© 2026 FIOR Group Limited
Visit our website: fior.group

FIOR and Multiverse Computing: a controlled architecture

A combined model and enforcement architecture

FIOR announced its partnership with Multiverse Computing on 11 August 2026, bringing its compression and deployment capabilities together with FIOR agent identity, policy enforcement, revocation and audit technology. The collaboration covers sovereign, edge, on-premises and regulated environments, with integration and deployment activity organised around customer requirements.

The operating proposition is straightforward. Multiverse makes a suitable model smaller and more economical to execute. FIOR controls the authority of the agents using that model and the systems they can reach. A locally running model may still issue a harmful instruction or call a tool beyond its remit; an independent enforcement layer can prevent an unauthorised action even where the generated plan appears plausible.

To meet the nationally bounded requirement, the complete customer adaptation process must take place inside the approved UK boundary. An eligible base model and toolchain are imported, checked and retained. Customer data preparation, teacher inference, synthetic data generation, distillation, compression healing, fine-tuning and evaluation then run on approved local infrastructure. The resulting weights, adapters, checkpoints and evaluation artefacts remain there.

A generic teacher API outside that boundary cannot be used with protected customer material. In the sovereign architecture described here, all customer data and all customer-specific training and derived model artefacts remain inside the approved jurisdiction from their creation. Where the base model was originally trained overseas, that historical pretraining remains distinct from the nationally bounded adaptation and deployment process. Exclusively domestically originated pretraining is a separate procurement requirement with materially different economics.

Lifecycle stage	Required UK boundary	Acceptance evidence
Import and provenance	Approved model and software copies	Hashes signatures licences and component inventory
Training and optimisation	UK compute storage and administrators	Job records access logs and blocked outbound routes
Derived artefacts	UK checkpoints adapters logits and caches	Inventory backup locations and deletion controls
Inference and retrieval	UK model endpoints and authorised records	Traffic inspection and retrieval permission tests
Agent actions	Customer-controlled FIOR enforcement points	Denied action tests delegation limits and approval records
Support and updates	Customer-approved signed packages and access	Offline operation test rollback and support procedure

© 2026 FIOR Group Limited
Visit our website: fior.group

Lifecycle stage	Required UK boundary	Acceptance evidence
Exit and recovery	Retained weights configurations and local evidence	Restore exercise licence review and substitution trial

Conditions needed for full operational sovereignty

Our sovereign deployment design combines locally resident model adaptation and inference with customer control of keys, administrators, updates, agent permissions and audit retention. The customer should be able to retain and operate the delivered artefacts through a supplier outage, subject to clearly agreed licence rights. No mandatory vendor cloud control plane should be necessary for routine operation.

Our approach builds sovereignty through operational control, supported by appropriate contracts and continuity arrangements. Multiverse is a Spanish supplier, and our partnership demonstrates how overseas expertise can contribute to a locally controlled service. We recommend explicit provisions covering remote access, intellectual property, licensing and software continuity so that the customer can sustain the service and manage remaining dependencies.

Sovereign agentic security

Our role is to make human and institutional authority enforceable when software agents act at machine speed. We give agents a cryptographic identity, bind them to delegated permissions and apply policy checks before they access protected tools, APIs, workflows or data. Higher-impact actions can require human approval; local evidence records what was permitted or denied.

For example, an engineering assistant could read approved maintenance instructions and draft a work order while being denied permission to alter a live control system. A public-sector casework agent could retrieve records within its remit without acquiring general access to an entire citizen database. These are illustrative policies, not claims that a model can establish legal authority on its own.

Effective implementation requires coverage of every consequential route. If an agent can bypass the gateway through a direct database connection, the policy boundary is incomplete. Tool credentials, network paths, privilege escalation and human-approved exceptions must therefore be included in testing. NCSC guidance supports secure infrastructure and signed model artefacts; the government AI cyber security code also addresses protection of training data, weights and software supply chains. [10][11]

© 2026 FIOR Group Limited
Visit our website: fior.group

Offline security has an unavoidable timing constraint. A disconnected device cannot instantly learn about a revocation made elsewhere. Local revocation, short-lived permissions, bounded delegation and preloaded policy can limit exposure, followed by synchronisation when a trusted connection becomes available.

Deployment in cloud, on premises and portable environments

Customer controlled cloud

A approved cloud deployment can pool demand, share administration and scale capacity across several institutions. Our software can run within that environment under customer operation, while the model and retrieval services remain inside the same approved perimeter. Software deployed on rented infrastructure does not have to become a supplier-operated SaaS service.

The cloud's underlying management plane still matters. Buyers should assess who can administer hosts, access support dumps, control encryption keys or suspend service. UK location does not settle overseas jurisdiction or administrator access.

Dedicated or appropriately isolated tenancy, customer-held keys, explicit support arrangements and tested backups improve control, but legal and operational dependencies must be assessed separately.

On premises and air gapped operation

On-premises deployment is attractive where organisations already maintain secure facilities, have stable demand or must integrate with isolated operational networks. It permits local custody of weights, prompts, identity policy and logs. Compression can reduce the number and class of machines required, easing deployment into an existing estate.

The customer assumes responsibility for capacity, patching, monitoring, physical security, recovery and incident response. Two sites, spare hardware and an offline update process may be necessary for critical services. Air-gapping can limit external exposure, but it also complicates vulnerability management and does not prevent malicious insiders, contaminated media or unsafe model outputs.

Lightweight portable systems

Portable AI can support field engineers, emergency responders, inspectors and deployed personnel when a remote service is unavailable or unsuitable. A compact model can retrieve local manuals, summarise observations or draft structured

© 2026 FIOR Group Limited
Visit our website: fior.group

reports. Our local operation carries identity and authority controls into the field without requiring a connection to FIOR.

The strategic advantage is continuity at the point of use. Data can be processed where it is collected and an authorised workflow can continue through communications outages. A portable assistant can also avoid broadcasting sensitive operational patterns to a distant model service. These benefits depend on the complete application remaining local, including embeddings, speech processing and retrieval where used.

“Lightweight” and “man-portable” require measurements for the selected configuration. Procurement should specify total carried mass, battery endurance, thermal limits, time to first token, output speed, context length, concurrent tasks and environmental tolerances. Model loading, storage encryption and audit overhead must be included. We specify device performance and carried weight for the selected model, hardware and mission profile rather than assign a universal figure to every portable deployment.

Physical capture creates risks that a data centre deployment does not face. Secure boot, encrypted storage, protected keys, controlled removable media and limited local privileges should accompany agent enforcement. Human authority can remain bounded through preapproved permissions, but destruction or compromise of hardware must be part of the threat model. A portable system is useful for a defined mission; it is not a replacement for all cloud capability.

The significance of no FIOR SaaS or telemetry

We provide customer-operated software without a FIOR SaaS dependency and receive no customer data or telemetry. Prompts, outputs, agent activity and diagnostic metadata do not need to flow to FIOR. Policy enforcement and operational records remain within the customer’s environment, preserving direct control over information that may disclose sensitive operational intentions.

We recommend applying the same sovereignty boundary to the complete deployment lifecycle, including installation, licence management, updates, crash handling and support. Supporting infrastructure such as identity services, registries and third-party libraries should remain within the approved boundary. Operational acceptance should demonstrate the agreed functions with external connectivity blocked. Any support material shared by the customer should be minimised and customer-authorised.

© 2026 FIOR Group Limited
Visit our website: fior.group

No telemetry to FIOR must coexist with effective customer monitoring. Local security and audit logs remain essential, with access restrictions, retention policies and protected export to a customer-controlled security operations centre. The advantage is customer custody of evidence. It does not justify running without evidence or assuming the rest of the deployment has no data processors.

Total cost, capacity and quality economics

Compare services at the same acceptance threshold

A meaningful comparison holds task quality, throughput, availability, data protection and security requirements constant. It includes training or adaptation, model licensing, infrastructure, staff, support, monitoring, retrieval, integration, human review and exit. A price per token measures only part of a service, and different tokenisers or reasoning patterns can produce different token counts for the same task.

The following scenario is illustrative, not a FIOR–Multiverse quotation or an independently measured benchmark.

It compares a compact customer-operated system with a representative hosted model API and a larger model on dedicated approved cloud compute. The latter distinction is important: using an external API and renting cloud hardware to run your own weights involve different economics and control rights.

Workload and commercial assumptions

Assume 10 million accepted tasks per year, each consuming 3,000 input tokens and 1,000 output tokens, without retries. That is 30 billion input and 10 billion output tokens annually. Quality and service-level parity are assumptions to be tested. A three-year undiscounted period is used; VAT, financing and corporation tax are excluded. All local costs are planning allowances.

For a transparent API reference, Anthropic lists Sonnet 4.5 at US$3 per million input tokens and US$15 per million output tokens. This is a named comparison point, not a claim that it is the cheapest or best current service. [19] At an assumed US$1 = £0.80, the workload costs £192,000 annually in model calls: (30,000 × $3 + 10,000 × $15) × 0.80. The exchange rate is a modelling assumption, not a current market quote.

For the local compact option, assume two nodes cost £60,000 combined, adaptation and validation cost £60,000 initially, and integration costs £40,000 initially. Annual allowances are £20,000 for model rights and support, £60,000 for operations, £8,000 for power and hosting and £5,000 for hardware maintenance. We assume two FIOR

© 2026 FIOR Group Limited
Visit our website: fior.group

gateways at our US$15,000 annual price per gateway, using the same exchange assumption: £24,000 annually. Actual gateway count and scope require a deployment quote.

The dedicated large-model cloud option assumes a resilient eight-accelerator estate, including redundancy within that count, at £2.50 per accelerator-hour continuously: 8 × 8,760 × £2.50 = £175,200 annually. That hourly rate and capacity are illustrative, not a provider quotation or validated architecture. Add £60,000 operations, £30,000 storage and network, £20,000 model rights and support and £24,000 FIOR annually, plus £60,000 initial adaptation and £40,000 integration.

The API option includes £60,000 initial integration and assurance, £40,000 annual operations, £20,000 annual retrieval and logging infrastructure and £24,000 FIOR annually. Including FIOR in all three options isolates the model and hosting choice. External provider retention, region and terms would still require separate approval; local FIOR enforcement does not make an external model endpoint sovereign.

Under these assumptions, compact local operation costs approximately 42% less than the API option and 50% less than the dedicated large-model cloud option over three years. These are scenario outputs, not general savings claims. No allowance is added for avoided breaches, productivity gains or sovereign economic multipliers because those benefits have not been measured.

Capacity is an investment gate

Ten million tasks per year average approximately 0.317 tasks per second across continuous operation. At 1,000 output tokens per task, the system must sustain about 317 output tokens per second on average, while processing about 951 input tokens per second. Business-hour concentration, latency targets and bursts raise the required peak capacity considerably.

The assumed two-node compact system must meet that demand during the agreed failure condition. If it cannot, additional nodes and licences change the result. The dedicated cloud estate must be tested on the same basis. Neither machine count is a performance assertion. A pilot should measure prefilling, output generation, retrieval and policy overhead together rather than extrapolate from a single-user demonstration.

© 2026 FIOR Group Limited
Visit our website: fior.group

Volume and price sensitivity

The table below holds compact capacity and fixed costs constant within a validated capacity range. API costs vary with task volume. Below that range, buyers should consider a smaller local system; above it, they must add capacity. The API column includes the same £60,000 initial and £84,000 annual non-token costs as the base case.

The base API model-call cost is £0.0192 per task. With the stated fixed costs, three-year break-even occurs at approximately 3.45 million tasks annually: (£511,000 - £312,000) ÷ (3 × £0.0192). This is a planning threshold for this configuration, not a universal threshold for sovereign AI.

A 50% reduction in API model-call spending reduces its base-case three-year TCO to £600,000. A 75% reduction makes it £456,000, cheaper than the assumed compact system. Prompt caching, batching, smaller hosted models and negotiated rates can therefore overturn the financial result. Published API pricing includes alternative charging mechanisms; buyers must model their eligible workload rather than assume list rates apply to every token. [19]

Local overhead can reverse the result too. An additional £50,000 of initial adaptation and £50,000 annually for staffing or capacity adds £200,000 over three years, making compact TCO £711,000. Dedicated approved cloud costs fall if accelerators can be scheduled for fewer hours or shared effectively; they rise if throughput, redundancy or geographic resilience require a larger estate. The comparison is sensitive to utilisation and service design.

Quality and human review can outweigh compute

The correct denominator is accepted tasks completed within the approved control boundary. Retries, rejected outputs, escalations and human corrections must be included. If the compact system requires one extra minute of human review per task on just 5% of 10 million annual tasks, at £30 per staff-hour, the added annual cost is £250,000. That exceeds the base-case annualised saving against the API option of about £125,667.

A route that sends 10% of the base workload to the reference API would add £19,200 annually in model calls before extra integration and assurance. Such escalation can improve quality economically, but it is incompatible with a strict no-external-processing requirement for protected material. Sensitive escalations may need a larger locally hosted model or a human reviewer instead.

© 2026 FIOR Group Limited
Visit our website: fior.group

Sovereignty benefits

Some benefits are constraints rather than cash savings. A prohibited external transfer or a mandatory offline mission can make an external API infeasible regardless of its price. In that case, compare the compact system with another compliant local or sovereign-cloud design. It would be misleading to claim savings against an option the customer could never lawfully or operationally adopt.

Where both alternatives are feasible, quantify downtime exposure and recovery time separately from compute savings. Use the organisation's incident data and service criticality to estimate avoided loss. National value retention should likewise be measured through actual UK employment, skills, intellectual property and procurement spending; the entire contract value does not automatically remain in the country.

Energy benefits should be measured as watt-hours per accepted task at representative load, including idle time and shared infrastructure allocation. Lower model memory can enable lower-power devices, but a lightly used local server can consume more energy per useful task than a highly utilised hosted service.

No percentage energy reduction can be inferred directly from a percentage reduction in parameters.

Procurement and implementation priorities

Start with a bounded service

A credible first procurement should select a sensitive, repeatable workload with an accountable owner and a measurable existing process. Define data classifications, acceptable failure modes, offline duration, peak demand and authorised actions before selecting a model. Decide whether UK residency, domestic administration, continuing licence rights or exclusively domestically originated technology is actually required.

The evaluation should include a compact open-weight baseline without proprietary compression, a Multiverse-optimised candidate, and a more capable reference model where permissible. This establishes the incremental value of compression rather than comparing only against an oversized alternative. The same retrieval corpus, quality rubric, security controls and human-review accounting should apply across candidates.

© 2026 FIOR Group Limited
Visit our website: fior.group

Require a sovereignty evidence pack

The supplier should provide the complete data flow map, model provenance and licences, software component inventory, administrator locations, key management design, support arrangements and continuity rights. Evidence should cover customer data and derived artefacts, including teacher outputs, cached logits, embeddings, backups and crash records. “No training on your data” is a narrower assurance than “your data is never processed outside the member state”.

Technical acceptance should demonstrate blocked external egress, local inference, local policy decisions, local audit capture, supplier-disconnection operation and recovery from backup. Security testing should attempt unauthorised tool actions, credential misuse, gateway bypass, stale delegation, policy rollback and tampered model loading. Test model quality after compression and again after material updates.

Portable acceptance should add battery and thermal testing, physical loss scenarios and operation through the full specified disconnection period. Cloud acceptance should add management-plane and support-access assessment. On-premises acceptance should verify restoration without supplier intervention.

All deployment forms need a tested way to replace a model without rebuilding the entire application.

Allocate decisions to the right stakeholders

Boards and accounting officers should approve risk appetite, value and continuity requirements. Service owners should own outcomes and human escalation. Security teams should own the enforcement boundary, incident handling and technical acceptance. Data protection and legal teams should establish lawful processing, transfer arrangements and model rights. Commercial teams should secure exit, support and continuing operation terms.

Government funders can make efficient sovereign deployment easier by supporting shared evaluation datasets, locally resident adaptation facilities and reusable assurance patterns. Sector bodies can develop benchmarks reflecting actual UK services. Procurement should reward demonstrated interoperability and tested operational independence so that sovereign investment does not simply replace one form of lock-in with another.

An indicative pilot can be completed in stages over approximately 90 days: establish the baseline and boundary, build and evaluate model candidates, integrate

© 2026 FIOR Group Limited
Visit our website: fior.group

enforcement, then conduct operational and exit tests. This is a planning sequence, not a delivery commitment. Rollout should follow evidence of task quality, service capacity and authorised behaviour, with price and capacity assumptions refreshed before purchase.

© 2026 FIOR Group Limited
Visit our website: fior.group

9 | Priority sectors and adoption pathways

The case for locally controlled AI is strongest in defence-support systems, justice and public administration, financial services, health and life sciences, transport, communications, energy, emergency response and commercially sensitive research. Legal and security requirements differ substantially by country. A UK-specific sectoral example illustrates design choices, but it should never be presented as a Commonwealth-wide statutory mandate.

Stakeholders	Workloads likely to value sovereign operation	Main reason for adoption
Defence intelligence and national security	Intelligence support technical manuals planning assistance	Classified data mission continuity and controlled authority
Central government devolved administrations and councils	Casework policy drafting secure records search	Citizen information institutional accountability and service resilience
Police justice and emergency services	Evidence review incident support field reporting	Sensitive investigations evidential integrity and disconnected use
Banks insurers payments and market infrastructure	Internal analysis investigations operations assistants	Confidential data third-party concentration and action control
NHS health research and care providers	Records summarisation administrative support research	Patient confidentiality provenance and clinical governance
Energy nuclear water transport and telecom operators	Maintenance incident response engineering knowledge	Continuity operational data protection and edge execution
Aerospace advanced manufacturing and defence suppliers	Design assistance technical search supplier workflows	Intellectual property export-sensitive material and local operation
Universities life sciences and professional services	Research analysis privileged document review	Valuable datasets confidential know-how and client obligations

Defence and emergency services may prioritise disconnected execution and protection of operational data. Governments must preserve lawful processing and contestable decisions. Financial firms need segregation of duties, third-party concentration controls and detailed records of action. Health organisations must test clinical risk, information governance and applicable device rules. Industrial and telecoms operators should isolate AI assistants from safety-critical controls unless independently justified and tested.

© 2026 FIOR Group Limited
Visit our website: fior.group

10 | A Commonwealth implementation programme

A voluntary 36-month programme could translate shared principles into deployable national capability. During months 0–6, member-led working groups would define risk classes, governance-mapping templates, sovereignty evidence packs and three representative use cases. During months 6–18, participating organisations would pilot bounded workloads in different jurisdictions, measure quality and cost, test offline operation, and run independent security and model evaluations. During months 18–36, the programme would publish reusable patterns, build assessor capacity, expand local technical teams and compare operational results across members.

Each national pilot should designate an accountable public or enterprise service owner and an independent evaluator; specify who holds keys, model weights and audit records; demonstrate an exit from a remote supplier; and test the failure modes relevant to the service. Shared technical evidence should be portable across member countries, while any regulatory approval remains a national decision.

The most useful joint output would be a practical Commonwealth Sovereign AI Deployment and Assurance Framework: common terminology, a checklist of national choices, governance-mapping patterns, reference controls, assurance test reports, and clear statements of what has and has not been independently assessed.

Conclusions

The Commonwealth can be a serious partnership for AI sovereignty without attempting to reproduce every element of the global frontier technology industry. Its members can cooperate on design, security tests, local skills, model adaptation and purchasing power while retaining control over their own data, institutions and laws. There are useful contributions from across the association, up and down the AI value chain.

This paper proposes a practical path: deploy models that are economical and suitable for the task; keep sensitive adaptation and operation within the required jurisdiction; map delegated authority before agents are released; enforce that authority at runtime; and assemble evidence that the system will remain secure and useful.

We propose three complementary workstreams: (i) practical sovereign deployment architectures using compact and, where needed, frontier models; (ii) governance mapping and enforceable agent identity, delegated authority and pre-action security controls; and (iii) shared design principles and independently evidenced assurance,

© 2026 FIOR Group Limited
Visit our website: fior.group

adapted to each member's laws and institutions. Cooperation should create interoperability, not a new source of dependency.

The operational standard should be simple to state and demanding to test: an AI service must perform its intended work, protect sensitive information, continue through defined disruption, and act only within authority that a human institution has lawfully delegated.



David Williams

Founder

FIOR Group



Anthony Halley

Director, Global Strategy & Programmes

Commonwealth Security & Resilience Group

© 2026 FIOR Group Limited
Visit our website: fior.group

References and source notes

This joint paper incorporates FIOR's sovereign-deployment technical analysis and its illustrative economic model. Statements about the FIOR CRT result are limited to the scope of its independent assessment. Economic assumptions are not forecasts or supplier quotations. UK-specific statutes and standards are examples and require local jurisdictional adaptation.

[1] Multiverse Computing. FIOR Group and Multiverse Computing announce strategic partnership to deliver secure edge and sovereign AI infrastructure. 11 August 2026.

https://multiversecomputing.com/resources/fior-group-and-multiverse-computing-announce-strategic-partnership-to-deliver-secure-edge-and

[2] Council for Science and Technology. AI sovereignty preserving UK agency in an AI enabled world. Published 19 October 2026; letter dated 30 July 2026.

https://www.gov.uk/government/publications/ai-sovereignty-preserving-uk-agency-in-an-ai-enabled-world/ai-sovereignty-preserving-uk-agency-in-an-ai-enabled-world-web-version

[3] Center for a New American Security. Sovereign AI Index. Dataset described as current to June 2026.

https://interactives.cnas.org/reports/sovereign-ai-index/

[4] Federal Government of Germany. Fragen und Antworten zu Künstlicher Intelligenz. Current policy overview, and Summit on European Digital Sovereignty commitments.

https://www.bundesregierung.de/breg-de/aktuelles/faq-kuenstliche-intelligenz-2458970

https://www.bundesregierung.de/breg-de/aktuelles/summit-on-european-digital-sovereignty-delivers-landmark-commitments-for-a-more-competitive-and-sovereign-europe-2394368

[5] Department for Science Innovation and Technology. AI firms pioneering drug discovery cheaper supercomputing and more get first backing through UK Sovereign AI. 16 April 2026.

https://www.gov.uk/government/news/ai-firms-pioneering-drug-discovery-cheaper-supercomputing-and-more-get-first-backing-through-uks-sovereign-ai

[6] Cottier B and colleagues. How much does it cost to train frontier AI models. Epoch AI. 3 June 2024. Historical estimates and conditional projection.

https://epoch.ai/publications/how-much-does-it-cost-to-train-frontier-ai-models

[7] Meta. Introducing Llama 3.1 Our most capable models to date. 2024.

https://ai.meta.com/blog/meta-llama-3-1/

[8] Stanford HAI. The 2025 AI Index Report. Research and Development overview.

https://hai.stanford.edu/ai-index/2025-ai-index-report/research-and-development

[9] Multiverse Computing. The CompactifAI pipeline. Undated methodology page as retrieved.

https://multiversecomputing.com/compactifai-pipeline

[10] National Cyber Security Centre. Guidelines for secure AI system development Secure deployment.

© 2026 FIOR Group Limited
Visit our website: fior.group

https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development/guidelines/secure-deployment

[11] Department for Science Innovation and Technology. Code of Practice for the Cyber Security of AI.

https://www.gov.uk/government/publications/ai-cyber-security-code-of-practice/code-of-practice-for-the-cyber-security-of-ai

[12] National Protective Security Authority. Critical National Infrastructure.

https://www.npsa.gov.uk/about-npsa/critical-national-infrastructure

[13] Ministry of Defence. The Strategic Defence Review 2025 Making Britain Safer secure at home strong abroad.

https://www.gov.uk/government/publications/the-strategic-defence-review-2025-making-britain-safer-secure-at-home-strong-abroad/the-strategic-defence-review-2025-making-britain-safer-secure-at-home-strong-abroad

[14] Cabinet Office. Government Security Classifications Policy.

https://www.gov.uk/government/publications/government-security-classifications/government-security-classifications-policy-html

[15] Bank of England and Financial Conduct Authority. Artificial intelligence in UK financial services 2024. 21 November 2024.

https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024

[16] NHS England Digital. Artificial intelligence. Information governance guidance overview. Updated 11 May 2026.

https://digital.nhs.uk/data-and-information/information-governance/guidance/artificial-intelligence

[17] NHS England Digital. NHS and social care data off shoring and the use of public cloud services. Updated 28 May 2025.

https://digital.nhs.uk/data-and-information/looking-after-information/data-security-and-information-governance/nhs-and-social-care-data-off-shoring-and-the-use-of-public-cloud-services

[18] Information Commissioner's Office. What is an international transfer of personal information.

https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/international-transfers/a-guide-to-international-transfers/what-is-an-international-transfer-of-personal-information/

[19] Anthropic. Claude Platform pricing. Sonnet 4.5 standard input and output list prices as retrieved. Prices and eligibility for alternative charging arrangements may change.

https://platform.claude.com/docs/en/about-claude/pricing

[20] Department for Science Innovation and Technology. Trusted third-party AI assurance roadmap. 3 September 2025.

https://www.gov.uk/government/publications/trusted-third-party-ai-assurance-roadmap/trusted-third-party-ai-assurance-roadmap

© 2026 FIOR Group Limited
Visit our website: fior.group

[21] AI Safety Institute. Approach to evaluations. Published evaluation methodology and limitations.

https://www.gov.uk/government/publications/ai-safety-institute-approach-to-evaluations/ai-safety-institute-approach-to-evaluations

[22] National Cyber Security Centre. Cyber Resilience Test Facilities Information for technology vendors. Assessment process and report ownership.

https://www.ncsc.gov.uk/schemes/cyber-resilience-test-facilities/information-for-vendors

[23] National Cyber Security Centre. One small step for Cyber Resilience Test Facilities one giant leap for technology assurance. Risk-based reporting and the absence of pass or fail certification.

https://www.ncsc.gov.uk/blog-post/one-small-step-for-cyber-resilience-test-facilities-one-giant-leap-for-technology-assurance

[CW1] Commonwealth Secretariat. Commonwealth membership, population and small-state profile. https://thecommonwealth.org/

[CW2] World Health Organization and Commonwealth Secretariat. Partnership release, 15 September 2026. https://www.who.int/news/item/15-09-2026-who-and-commonwealth-secretariat-renew-partnership-to-strengthen-health-systems-and-advance-universal-health-coverage

[CW3] Commonwealth Secretariat. State of the Commonwealth Report 2026: Resilience through Diversity, 2 November 2026. https://thecommonwealth.org/CHOGM2026/news/democratic-economic-and-environmental-resilience-integral-commonwealth

[CW4] UK Government. UK AI Hardware Plan, 8 June 2026. https://www.gov.uk/government/publications/uk-ai-hardware-plan

[CW5] UK Government. Sovereign AI and British Business Bank invest in Fractile, 11 November 2026. https://www.gov.uk/government/news/sovereign-ai-and-british-business-bank-invest-in-british-firm-developing-the-next-generation-of-ai-hardware

[CW6] AI Security Institute. About and evaluation research. https://www.aisi.gov.uk/about

[CW7] Commonwealth Research Group. AI Security Summit 2026, 24–25 November, Marlborough House, London. https://www.cwsrg.org/ai-securitysummit2026

© 2026 FIOR Group Limited
Visit our website: fior.group