

AI

Trusted AI for the Integrated Force

Accelerating Warfighting Readiness
through Assured AI, Data and
Decision Advantage

WHITEPAPER

Executive summary

The Defence Investment Plan 2026 establishes that true mission advantage will no longer be won by traditional platforms alone, but by the rapid, sovereign exploitation of data, software, artificial intelligence and autonomy. The DIP sets out total Defence spending of £298bn over FY26/27 to FY29/30, with investment choices designed to move Defence to warfighting readiness, back British industry and strengthen UK leadership with allies.

Artificial Intelligence is central to this transformation. The DIP states that AI is the most important of the technologies changing how war is fought: a core enabler of military power and institutional effectiveness, accelerating decision-making, expanding sensing and effect, increasing accuracy and lethality, and changing how advantage is generated and sustained.

However, the principal Defence AI challenge is not the creation of ever more impressive models. It is the operationalisation of trusted AI at scale. Defence has seen many promising demonstrations, prototypes and laboratory capabilities. The next task is to convert AI into assured, deployable and continuously improving mission capability in environments where decisions matter, communications are contested, adversaries adapt and failure is not an option.

Thales recommends that Defence moves from a model-centric approach to a harness-centric approach. In mission-critical environments, the AI model is only one component of capability. The decisive advantage comes from the surrounding system: data, assurance, cyber security, governance, human oversight, monitoring, deployment architecture, operating procedures, feedback loops and

continuous improvement. This is the AI Harness. It is what turns AI from a promising algorithm into trusted operational capability.

The DIP provides the policy and investment framework to establish this at national scale. Specific AI investments such as Taskforce RAID, Project FRONTIER, Land AI C2 and AI for underwater dominance are important accelerants, but they cannot be the whole answer. AI must traverse the whole of Defence through a distributed, coordinated model that enables every Front Line Command, enabling function, acquisition route, industrial partner and operational community to adopt trusted AI against its own mission problems while reusing common standards, assurance patterns, data foundations and deployment architectures.

This is the critical distinction. Defence does not need a single AI programme sitting at the centre of MOD. It needs a federated AI ecosystem: coordinated enough to avoid fragmentation, duplication and incompatible standards; distributed enough to deliver at the edge of operational need; and open enough to harness the full breadth of UK industry, SMEs, academia, allies and investors.

Thales recommends that Defence establishes a Trusted AI Harness approach across the Integrated Force, focused on six priorities:

- **Embed assured AI into the Digital Targeting Web** and the wider Digital Backbone.
- **Adopt a distributed but coordinated AI delivery model** across MOD, industry and operational users.
- **Deploy frugal AI at the edge** in denied, degraded, disconnected and constrained environments.
- **Establish common assurance and governance** for agentic AI, autonomy and AI-enabled command support.
- **Use AI to transform Defence productivity** as well as frontline capability.
- **Anchor trusted AI in the UK industrial base** with sovereign capability, export potential and allied interoperability.

AI wherever it accelerates understanding, human judgement wherever accountability matters, integrated always.

Accelerating AI from prototype to operational capability

The transition to AI-enabled Defence changes the role of platforms, the skills required by personnel, the way industry contributes and the route by which capability reaches the front line. Platforms will remain essential, but their effectiveness will increasingly depend on the speed and trust with which they can sense, decide, connect and act across domains.

The DIP is clear that the Integrated Force must move from individual platform advantage to all-domain integration, enabled by data, autonomy and software-defined warfare. The most important Defence AI question is therefore not whether AI can work in a controlled environment. It is whether AI can be deployed, trusted, governed, secured and improved in operational conditions.



RECOMMENDATION

Establish a Defence AI Harness as the common route from

prototype to deployed capability.

Defence should require AI capabilities across central programmes, Front Line Commands, enabling organisations and industrial delivery routes to demonstrate not only model performance, but operational readiness. This should include data lineage, assurance evidence, cyber resilience, human-machine interface design, monitoring, auditability, drift detection, update pathways and clear accountability.

Move AI delivery from technology demonstrations to operational problem ownership.

AI delivery should be organised around priority Defence problems, with named operational sponsors, industrial partners and delivery teams responsible for moving from minimum viable capability to deployed and sustained operational use. This will reduce duplication, accelerate learning and ensure that AI investment is measured by mission effect rather than experiment volume.

A distributed but coordinated AI model for MOD

AI will not succeed in Defence if treated as a single central fund, a single specialist unit or a linear innovation pipeline. The operational demand for AI is highly distributed: in maritime autonomy, land recce-strike, air command and control, space domain awareness, cyber and electromagnetic operations, intelligence fusion, logistics, acquisition, engineering support and enterprise productivity.

The DIP already points to this distributed reality. AI is embedded in the Digital Targeting Web, Project ASGARD, the Maritime Fighting Web, Hybrid Navy, Collaborative Combat Aircraft, Integrated Air and Missile Defence, Defence-wide Secret Cloud, Project FRONTIER, Land AI C2, underwater dominance, Digital Warfighters and MOD enterprise automation. These must not become disconnected AI islands. They must form a trusted and coherent AI ecosystem.



RECOMMENDATION

Create a federated AI operating model across Defence. MOD should establish a common AI governance, assurance and deployment framework that allows AI solutions to be developed close to operational need while aligning to shared standards. The centre should set the rules of the road; the domains and enabling functions should own mission outcomes; and industry should provide reusable, assured capability components that can be adapted across contexts.

Use central initiatives as accelerants, not as the whole answer. Taskforce RAID and Project FRONTIER can provide pace, protected investment, specialist talent and early operational exemplars. But Defence advantage will come when those exemplars are re-used and scaled across the wider enterprise through common architectures, common assurance evidence, common data products and clear routes into enduring programmes.

Build an industry-wide AI capability mesh. Primes, SMEs, start-ups, universities and specialist suppliers should be orchestrated around open interfaces, reusable assurance artefacts, secure data access routes and shared test environments. This will allow Defence to exploit specialist innovation without creating fragile point solutions that cannot be sustained or integrated.

Building the AI Harness

A model alone does not create military capability. In mission-critical environments, operators need to know whether AI can be trusted, understood, deployed, secured, constrained, monitored and improved. These questions are answered by the harness around the model, not the model alone.

The AI Harness should become a sovereign Defence capability. It should provide the common architectural, assurance and governance layer required to deploy AI across classifications, networks, platforms, headquarters, tactical edge environments and coalition settings.

- **Trusted data foundations:** provenance, labelling, quality, operational relevance and access control.
- **Mission-specific assurance:** evaluation, red teaming, safety cases and operational acceptance criteria.
- **Security by design:** cyber hardening, adversarial robustness, digital identity and model protection.
- **Human-machine teaming:** explainability, operator confidence, training, trust calibration and escalation.
- **Operational monitoring:** drift detection, confidence scoring, degradation modes and incident reporting.
- **Continuous learning:** feedback from operations, retraining pathways and configuration control.
- **Federated deployment:** edge, platform, headquarters, cloud and coalition environments.

Mandate minimum viable assurance for operational AI. Defence should adopt a minimum viable assurance model that enables rapid fielding without bypassing trust. Assurance should be risk-based, mission-specific and continuously updated, not treated as a late-stage approvals hurdle.

Embedding AI into the Digital Targeting Web

The Digital Targeting Web is one of the most important Defence investments in the DIP. It will connect sensors, deciders and effectors across domains, creating speed and choice in how Defence identifies, deconflicts, degrades or destroys targets. The DIP sets out £1.8bn for the Digital Targeting Web and £5.5bn for the Digital Backbone over FY26/27 to FY29/30, with a combined Digital Backbone and Targeting Web category of £7.5bn.

This cannot simply become another network programme. It must become a trusted decision architecture. AI will be required to ingest data, fuse sensor feeds, prioritise information, recommend options, support targeting workflows and maintain confidence in contested cyber and electromagnetic environments.

RECOMMENDATION



Integrate an assured AI decision-support layer into the Digital

Targeting Web. Defence should avoid every domain procuring bespoke AI stacks. A common AI layer should provide reusable services for data fusion, confidence scoring, explainability, model governance, mission audit, operator feedback and coalition information sharing.

Create standardised sensor-to-decision AI nodes across the Integrated Force. Platform-agnostic AI nodes should translate high-volume raw data from sensors, uncrewed systems, space-based assets, cyber sources and intelligence feeds into compressed, actionable and assured operational insight. These nodes should feed the Digital Targeting Web and enable faster decision-making while preserving human command accountability.

Deploying trusted AI at the edge

The future operating environment will be distributed, contested and coalition-based. AI will need to operate on uncrewed systems, crewed platforms, tactical headquarters, secure clouds and coalition networks. In many cases it will need to function with constrained bandwidth, low power, limited compute, degraded sensors, intermittent connectivity and adversary interference.

This makes frugal AI a Defence necessity. AI for Defence must be designed not only for performance, but for deployability under operational constraint. It must work where the mission happens, not only where compute is abundant.

RECOMMENDATION



Prioritise frugal AI for denied and degraded environments.

Defence should establish edge-AI acceptance criteria for capabilities required to operate under low bandwidth, low power, intermittent connectivity, cyber attack, GPS denial and electromagnetic disruption. This should include local inference, model compression, graceful degradation, known failure modes and clear confidence indicators.

Develop sovereign edge-AI reference architectures by domain. Reusable deployment patterns should be created for maritime, land, air, space and cyber environments. These should define how AI is deployed, updated, monitored, secured and assured across platforms and tactical networks, giving industry a clear path to integration.

Agentic AI, autonomy and accountable command

The next phase of AI will not simply detect, classify or recommend. Agentic AI systems will reason, plan, use tools, interact with other systems and increasingly take action. Once AI systems begin to act, the assurance question changes fundamentally. It is no longer only whether a prediction is accurate. It is whether behaviour can be controlled, constrained, monitored and held accountable.

This is directly relevant to AI-enabled swarms, autonomous vessels, uncrewed ground vehicles, Collaborative Combat Aircraft, digital planning assistants, cyber tools and automated decision-support systems. The DIP recognises the importance of Dependable AI and Responsible Use of AI as central to maintaining the trust of Defence personnel, Government and society.

RECOMMENDATION



Establish a Defence Agentic AI Control Framework.

Defence should define common controls for AI systems that plan or act. These should include authority boundaries, permitted actions, tool access, escalation rules, human approval points, audit trails, sandboxing, behavioural testing, kill-switch mechanisms and mission-specific constraints.

Create joint regulator-industry-operator assurance teams.

Defence should establish mixed teams of operators, engineers, human factors specialists, cyber experts, legal advisers, regulators and industrial partners to assure AI against live operational problems. This would make assurance a delivery accelerator rather than a late-stage blocker.

Human-machine teaming as the decisive design principle

AI will not replace command responsibility. The most effective Defence AI will put humans in a better decision loop, reducing cognitive burden and improving the speed, confidence and quality of decisions.

Modern Defence is not suffering from a lack of information. It is facing a crisis of understanding caused by unprecedented volumes of ISR, sensor, cyber, intelligence and open-source data. AI must help operators find the signal in the noise, move from data to insight, and move from insight to decision before the adversary does.


RECOMMENDATION



Make human factors mandatory in Defence AI delivery.

Every operational AI programme should include human factors design from inception. This should cover explainability, alert fatigue, workflow integration, trust calibration, escalation pathways, operator training, mission rehearsal and post-mission review.

Establish AI mission rehearsal environments. The DIP highlights common synthetic environments, Digital Warfighters and AI-enabled decision support. Defence should build AI mission rehearsal environments where AI behaviours can be tested, challenged and explained before deployment. This will be essential for swarms, autonomy, targeting, operational planning and agentic systems.



Using AI to transform the Defence enterprise

AI is not only a frontline capability. It is also a major opportunity to improve productivity across Defence. The DIP states that Defence should automate at least 20% of HR, finance and commercial functions by July 2028. It also states that AI tools have already been embedded across enterprise systems for at least 20,000 users, with agentic AI capabilities undergoing trials across Defence.

This matters because warfighting readiness depends on institutional speed. Defence cannot move at the pace required if scarce military, civilian and industrial talent is consumed by avoidable administrative burden.

RECOMMENDATION



Create a Defence AI Productivity Accelerator.

Defence should identify high-volume, high-friction workflows where AI can reduce burden and generate measurable productivity gains. Priority areas should include acquisition documentation, engineering knowledge retrieval, support contracts, logistics planning, lessons capture, HR casework, finance processing and commercial analysis.

Use AI to accelerate acquisition reform. The DIP introduces procurement segmentation, accelerated commercial pathways and problem-led industrial engagement through the NAD Group. AI should support demand signal analysis, supplier discovery, requirements traceability, evidence management, compliance checking and portfolio decision-making.



Establishing the UK trusted AI defence industrial base

Delivering trusted AI at scale will require a stronger UK AI defence industrial base. The DIP positions Defence as an engine for growth, with investment in R&D, frontier industries, skills, regional clusters and sovereign capability. It also commits £1.6bn to UK Defence Innovation over four years and sets a target for at least 10% of the annual Equipment Programme budget to be spent on novel technologies including AI, autonomy, quantum and directed energy weapons.

The novel technology profile shows the scale of the opportunity: AI and autonomy spend rises from £380m in FY26/27 to £1.3bn in FY29/30. But if this investment is fragmented, Defence will buy isolated functions rather than enduring advantage. AI is therefore both an operational priority and an industrial opportunity. The UK can build sovereign capability, create high-value skills, support SMEs and establish exportable trusted AI architectures for allies.

RECOMMENDATION



Create a Trusted AI Industrial Framework. Defence should establish a framework through which primes, SMEs, start-ups, academia and investors can contribute to a common AI ecosystem. This should include open architectures, reusable assurance patterns, synthetic test environments, secure data sandboxes, accreditation pathways and commercial routes from prototype to programme.

Adopt operational problem books for AI. Defence should publish challenge-led AI problem statements through appropriate secure channels, allowing industry, academia and SMEs to propose new approaches against real operational needs. This supports the DIP's emphasis on problem-led engagement, innovation and accelerated commercial pathways.

Design trusted AI for allied interoperability from the start. AI will be central to AUKUS, GCAP, NATO interoperability, the Digital Targeting Web, the Hybrid Navy, Collaborative Combat Aircraft and coalition operations. Trusted AI architectures should therefore be designed with data-sharing controls, assurance evidence, sovereign authority, exportability and allied integration built in from the outset.

Priority actions

1. Establish the Defence AI Harness. Define the common assurance, governance, security, data and deployment architecture required to move AI from prototype to mission capability.
2. Create a federated AI operating model across MOD. Coordinate AI standards centrally while enabling distributed delivery by Front Line Commands, enabling functions, industry and operational users.
3. Embed trusted AI into the Digital Targeting Web. Build a common AI decision-support layer for data fusion, confidence scoring, explainability, audit and operator feedback.
4. Prioritise AI at the edge. Develop frugal AI capabilities that operate in denied, degraded, disconnected and constrained environments.
5. Create a practical control framework for agentic AI. Define authority boundaries, escalation rules, behavioural constraints, monitoring and accountability for AI systems that plan or act.
6. Make human-machine teaming central to AI delivery. Embed human factors, mission rehearsal, training and trust calibration into every operational AI programme.
7. Use AI to accelerate Defence productivity. Apply AI to acquisition, support, logistics, HR, finance and commercial functions to release capacity for the front line.
8. Build a UK trusted AI industrial base. Give industry, SMEs and academia clear standards, problem statements, assurance pathways and routes to scale.

Summary

The Defence Investment Plan 2026 creates a decisive opportunity for the United Kingdom. It recognises that AI, autonomy, data and software are now central to operational advantage, lethality, productivity and industrial growth.

The risk is fragmentation: isolated pilots, domain-specific models and bespoke systems that demonstrate promise but fail to scale. The opportunity is to build a trusted, sovereign and assured AI ecosystem that connects people, data, machines and missions across the Integrated Force.

For Defence, success will not be determined by the largest model or the most impressive demonstration. It will be determined by the ability to deploy trusted AI at scale, assure it, secure it, govern it, integrate it with operators and continuously improve it under operational conditions.

The core message is clear: trusted AI is not only about algorithms. It is about creating the harness that allows AI to be used when it matters most, and the distributed industrial ecosystem that allows Defence to apply it everywhere it is needed.



350 Longwater Avenue
Green Park
Reading
RG2 6GF
United Kingdom

+44 (0) 118 943 4500

>thalesgroup.com<

