



Trusted, Sovereign AI: The Strategic Imperative for National Security, Innovation and Talent Development

THALES
Building a future we can all trust

Trusted, Sovereign AI: The Strategic Imperative for National Security, Innovation and Talent Development

Vincent Gicquel, Chief AI Officer, Thales Emarat Technologies
and Gabriel Rangoni, Vice President Strategy & Marketing, Thales
June 2026

Executive Summary

Artificial Intelligence (AI) has become the decisive lever of national power. It shapes how nations defend their interests, stabilise critical infrastructure and grow their economies.

If the intelligence behind these systems is designed, trained and governed elsewhere, nations inherit risks they cannot easily control. They face opacity in algorithms, dependencies on foreign platforms and cloud providers, and constraints that can undermine both policy autonomy and operational continuity.

That is why AI Sovereignty matters. It puts decisions, compliance and resilience under domestic authority. More than a technology framework, it serves as a national capability built through trust, talent and transparent collaboration.

For the Middle East, the urgency is even more pronounced. The UAE and Saudi Arabia are pursuing some of the world's most ambitious AI and digital-economy strategies, underpinned by national visions such as We the UAE 2031 and Saudi Vision 2030.

These countries are rapidly scaling smart-government services, advanced defence ecosystems, aviation hubs and energy transitions as domains where data sensitivity, operational continuity and regulatory alignment demand sovereign control.

With large-scale investments in talent, R&D, and national compute capacities already under way, AI Sovereignty offers the framework to ensure that innovation is developed on national terms, that strategic sectors remain resilient, and that AI becomes a lasting source of regional security, prosperity and influence.

In essence, AI Sovereignty is the disciplined ability to design, operate and control data, models, infrastructure and operations according to domestic law, mission policy and societal values.

This whitepaper, **Trusted, Sovereign AI: The Strategic Imperative for National Security, Innovation and Talent Development**, sets out Thales's framework for ethical, explainable and cyber secure systems operated end to end under national oversight, from data generation and collection to model training, deployment and audit.

Thales's perspective is straightforward and actionable. Nations should decide what remains domestic, what is shared and how collaboration occurs, so that autonomy is preserved where it matters and interoperability is pursued where it benefits.

In this, human oversight is non-negotiable. In mission-critical environments, AI must augment human operators and commanders rather than displace them. Above all, trust must be evidenced and demonstrated through data provenance, audit trails, explainability reports, robustness testing and certifiable security controls.

Thales's Trusted, AI Sovereignty framework rests on four pillars: **Trust, Sovereignty, Innovation and Future Readiness.**




Figure 1: Thales's Trusted, AI Sovereignty framework

With AI Sovereignty, national decisions stay within the nation's realm. Every decision is explainable, auditable and defensible. Talent pipelines, local IP and co development further contribute to national capability. Frugal design reduces size, weight and power (SWaP) and lowers energy footprint while delivering real time performance.

This white paper sets an action agenda: establish national AI assurance with mandatory explainability; operationalise data sovereignty end to end; adopt dual sovereignty that combines sovereign models for sensitive workloads and vetted open-source stacks for non-critical domains; prioritise frugal AI for continuity and sustainability; and institutionalise talent pathways.

With the right governance, architectures and talent, countries can convert AI dependency into a national capability with durable advantage.

1. Introduction: Why AI Sovereignty Matters

1.1 GLOBAL LANDSCAPE: GEOPOLITICAL SHIFTS, TECHNOLOGICAL DEPENDENCIES, SUPPLY CHAIN FRAGMENTATION

Today, the balance of national power is increasingly drawn by technology. Nations operate under continuous pressure from rapid innovation, fragmented supply chains and contested information environments.

AI is embedded across defence, civil aviation, maritime domains, energy systems and public services. The argument is no longer whether to adopt AI. It is about who governs the intelligence behind these systems, and on what terms.

AI Sovereignty answers that question with genuine national capability. It extends control across the entire lifecycle so that the intelligence, which helps build security, shape economic policy and strengthen public services to be designed, deployed and audited under national authority.

It replaces import led dependency with verifiable trust: explainable decisions, certifiable security, traceable provenance, and policy responsive data pipelines that keep systems aligned with national rules and obligations.

Three principles guide this approach:

Thales brings a dual DNA to AI Sovereignty that matches these demands.

Our defence grade trust reflects decades of rigorous engineering, certification readiness, cyber security, strict access controls and resilience in all conditions. These are the essential qualities needed to keep systems dependable even in the most challenging situations.

Our innovation delivers user centric design, software agility, cloud and UX know how, data analytics and rapid iteration. Combined, they enable national control without compromising speed, ensuring that the systems are explainable and auditable, with robust field capabilities, and adaptable to national civilian contexts and market needs.

1.2 DEFINING AI SOVEREIGNTY

An effective way to understand AI Sovereignty is to think of it as a national power grid: a distributed network of intelligence that operates independently yet remains connected under one trusted system. Each node functions securely on its own, ensuring that critical services continue even in moments of disruption. This model prevents dependency "blackouts," giving nations the confidence that their AI systems will remain transparent, resilient, and aligned with national priorities.

True sovereignty lies in maintaining control over how innovation is developed, shared, and applied.




Figure 2: The Three Principles

Nations can engage in open collaboration with startups, academia, and international partners while ensuring that sensitive data, infrastructure, and operations remain under their authority. This balance between openness and control defines the next phase of technological independence.

Thales supports this trajectory through co-development, working with public agencies, industry, and academia to localise engineering, transfer know-how, and align capability with national priorities.

A prime example is Thales Emarat Technologies in the UAE, which serves as a platform for joint research and technology transfer that nurtures local talent while applying global engineering standards. The emphasis is on building people, infrastructure, and governance, ensuring that critical expertise and intellectual property are retained inside the country, forming the foundation for durable sovereignty.

1.3 BUILDING AND GOVERNING AI ON NATIONAL TERMS

In the Middle East, AI Sovereignty is a strategic choice accelerated by national ambition. The United Arab Emirates has set out a clear vision through its National AI Strategy, grounded in the belief that data, not oil, will shape the future.

Saudi Arabia, through its \$40-billion AI investment plan, has similarly positioned AI at the heart of its transformation agenda. These programmes, aligned with national transformation agendas such as We the UAE 2031 and Saudi Vision 2030, signal the region's shared commitment to technological sovereignty, economic diversification, and knowledge-based growth.

These initiatives reflect a broader regional vision: to build domestic capability, develop indigenous talent, and govern technology on national terms.

AI Sovereignty, in this context, defines the authority to decide how innovation is developed, shared, and applied, under transparent frameworks for intellectual property, data access, and export-controlled technology transfer.

1.4 DATA SOVEREIGNTY: THE FOUNDATION OF AI SOVEREIGNTY

AI Sovereignty depends on an assurance stack that operators can rely on, which converts policy into practice and practice into evidence. Nations need clarity on roles, enforceable governance over data and models, and pipelines that integrate compliance into engineering rather than bolt it on after deployment.

Policy begins with a national authority that sets standards, risk classes, and compliance thresholds, defining rules for collaboration and export such as intellectual property, access rights, data sharing, and technology transfer. Operators own mission rules and approval frameworks, ensuring human-in-the-loop decision-making, while independent assurance teams audit pipelines and certify conformance with traceable evidence.

All of this is underpinned by sovereign data management. Nations must maintain end-to-end control over the entire data lifecycle, from generation and collection to storage, labelling, training, deployment, and auditing.

Data sovereignty ensures that information remains subject to domestic law and governance, underpinning the legality, safety, and explainability of AI systems. When sensitive data crosses jurisdictions or third-party clouds, nations risk losing control over model behaviour and compliance, eroding both trust and accountability.

A truly sovereign dataset is one that a nation fully governs throughout its lifecycle through:




Figure 3: The Lifestyle

At Thales, these principles are built into the AI development pipeline. Through federated learning environments like SaferLearn, data never leaves national boundaries, ensuring that sovereignty is maintained without sacrificing innovation, interoperability, or authority.

AI Sovereignty extends this foundation across the entire lifecycle, giving governments the ability to define, operate, and certify AI according to national priorities and security needs. In doing so, data sovereignty becomes the strategic national asset that fuels innovation, protects critical systems, and keeps AI decisions accountable to the people and institutions they serve.

1.5 CORTAIX, THALES'S ANSWER TO AI SOVEREIGNTY

Thales's approach to AI Sovereignty builds the framework that brings data sovereignty to life. We see the journey as a collaborative endeavour shaped by trust, transparency and shared innovation.

We work with governments, industries, and research partners to design AI systems that are ethical, explainable and secure by design, ensuring they reflect each nation's strategic priorities and regulatory frameworks.

At the heart of this approach is **CortAix**, Thales's global AI accelerator and innovation engine for mission-critical environments. CortAix brings together world-leading expertise in research, sensors and systems integration to deliver AI that is trusted, frugal and resilient under constraints.

Thales is a recognised leader in trusted, frugal, cybersecure, transparent and ethical AI, serving armed forces, aircraft manufacturers and operators of critical infrastructures.

The Group leads the European ranking of AI patent filers and, through the establishment of CortAix in 2024, unites over 800 engineers and PhD researchers across Canada, France, the UK, and Singapore.




Figure 4: Engineers and PhD Researchers Across the Globe

As Thales aims to extend the CortAlx ecosystem to key strategic regions, including the Middle East, the initiative is evolving to serve as a bridge between global research excellence and national co-development. The planned regional presence will strengthen local talent pipelines, accelerate sovereign innovation and ensure that the UAE and wider region benefit directly from the Group's trusted, mission-ready AI frameworks.

Thales's AI capabilities are embedded in more than 100 operational solutions, spanning information gathering, situational analysis, secure communications, maritime mine countermeasures and flight path optimisation. What distinguishes Thales is its hybrid model of AI: systems that deliver transparency, cybersecurity and ethical assurance without the excessive data and energy demands typical of large commercial models.

This **'magic square'** approach of balancing performance, frugality, security and ethics, positions Thales uniquely among global technology players. It exemplifies how trust serves as a verifiable advantage in building responsible, sovereign AI.

Operating under Thales's **TrUE AI** framework of Trusted, Understandable and Explainable AI, every CortAlx solution is built to ensure transparency, accountability and ethical integrity from design to deployment.

Delivering AI Sovereignty demands a structured framework that balances trust, control, innovation and adaptability. Thales defines this framework through four foundational pillars that guide its strategy:

- **Trust:** AI must be explainable, secure, and under human oversight. Decisions are auditable, traceable and defensible, ensuring confidence for both operators and regulators.
- **Sovereignty:** Nations retain control over data, models, pipelines and infrastructure across the entire AI lifecycle, maintaining independence from external systems or opaque algorithms.
- **Innovation:** Open, collaborative ecosystems, bringing together startups, SMEs and academic partners—thrive within clear governance structures that preserve national interests while encouraging creativity.
- **Future Readiness:** AI must remain adaptable to technological change, regulatory evolution and operational disruption. Flexibility and scalability are essential to ensure continuity even in contested or disconnected environments.

These pillars define how Thales enables nations in building trusted, sovereign and future-ready AI ecosystems that combine defence-grade trust with civil innovation, enabling autonomy without isolation. Through this framework, Thales ensures that AI is sovereign by design, aligned with national objectives and governed with integrity from the ground up.

2. Pillar One – Trust: The Foundation of Responsible AI

Trust is the foundation of any AI Sovereignty framework. At Thales, trust is built into every layer of design, deployment and decision-making. With AI Sovereignty, trust defines whether systems can be relied upon to act safely, transparently, and in alignment with human and national intent.

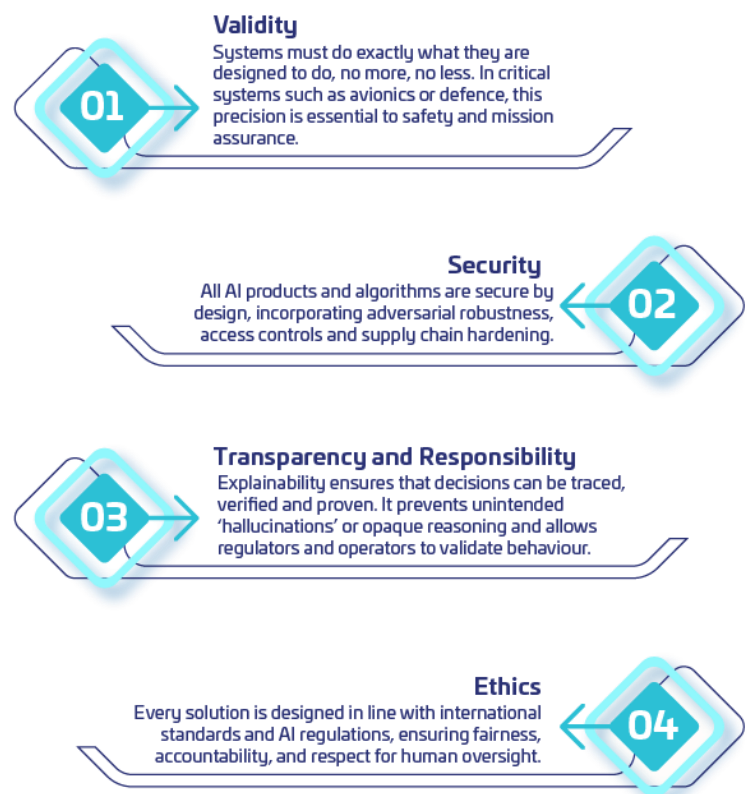
When systems are trusted, adoption accelerates, allowing governments, industries, and citizens to embrace AI with confidence and realise its full social and economic value.

2.1 DEFINING TRUSTED AI

Thales operationalises trust through its TrUE AI (Trusted, Understandable, and Explainable AI) framework. This ensures that AI systems remain transparent, secure and accountable from design to deployment.

Within CortAlx, TrUE AI principles are embedded as non-negotiable design standards. Every developer follows specific training and rigorous processes that integrate explainability, security, and ethical assurance into the entire lifecycle, from model conception to operational use.

In Thales's view, Trusted AI is built around four key principles that strengthen confidence in automation:



Thales aligns its AI assurance methodologies with emerging international frameworks such as ISO/IEC 42001, the EU AI Act, and OECD AI Principles, ensuring global compatibility while maintaining sovereign oversight.

2.2 EMBEDDING TRUST IN PRACTICE

In CortAIx, explainability and security are treated as non-negotiable design principles. This means every AI decision is explainable, every model verifiable and every output protected against tampering or misuse.

Through mechanisms such as watermarking, model validation, and cyber-resilience testing, Thales ensures that trust is measurable, demonstrable and enforceable.

This combination of transparency and security forms the core of what Thales calls 'Trusted AI Sovereignty' that nations can rely on for both operational performance and moral assurance. It ensures that humans remain firmly at the centre of decision-making.

3. Pillar Two – Sovereignty: Enabling Control and Autonomy

Sovereignty is the ability to define, deploy and govern AI on a nation's own terms. It ensures that intelligence, infrastructure and decision-making remain within national reach, which is critical for both security and self-determination.

At a time when data and algorithms move faster than policy, AI Sovereignty empowers countries to maintain control over their digital futures while engaging responsibly with the global ecosystem.

3.1 THE STRATEGIC VALUE OF AI SOVEREIGNTY

AI Sovereignty strengthens three essential pillars of national power: security, economic resilience and technological leadership.

By building domestic AI ecosystems, nations can overcome foreign dependencies that could potentially compromise operations or expose sensitive data. They also foster economic self-reliance, developing industries, technologies, and skills that stimulate long-term innovation and sustainable growth.

AI Sovereignty also stimulates homegrown innovation and high-value employment, keeping intellectual property, skills, and profits within national borders.

From a policy standpoint, AI Sovereignty allows governments to enforce their own regulatory, ethical and security frameworks rather than rely on imported systems or models that may not align with local

priorities. This is critical in sectors where AI decisions directly influence national stability and human safety such as defence, aerospace, energy, transportation and finance.

In this sense, sovereignty is an assertion of national will. It determines who holds authority over decisions that affect citizens, economies and critical infrastructure, and who can be trusted to uphold that authority when systems become increasingly autonomous.



What Makes Thales's Sovereign AI Different?

Thales combines defence-grade trust, edge- first design and human-in-the-loop assurance to make AI dependable under any conditions. Its architectures are secure, frugal and hardware-agnostic, ensuring continuity even when networks are disrupted. With transparency and human oversight built in, Thales turns Sovereign AI into a verifiable national capability: trusted, explainable and ready for mission use.

3.2 EDGE AI AS AN ENABLER OF SOVEREIGNTY

AI Sovereignty is about who controls data, models and infrastructure; edge AI is about where computation happens. Edge AI is one architectural pattern that can reinforce sovereignty, especially when continuity, latency and data-residency matter, but sovereignty can also be achieved with sovereign clouds and on-prem data centres where the mission allows.

Sovereignty is most powerful when it moves beyond theory into practice: when intelligence operates at the edge. Here, "edge" means placing AI inference (and, when appropriate, training/finetuning) close to sensors and users so that national authorities retain operational control even when central links are constrained.

Edge AI pushes intelligence and decision support directly onto platforms and sensors, whether in aircraft pods, naval systems, or ground units, allowing nations to maintain autonomy even under jamming, degraded

links, or disconnected environments. This supports sovereignty in three ways: (1) continuity under constraint (decisions proceed without reliance on external networks), (2) jurisdictional assurance (sensitive data can stay on national assets), and (3) safety and latency (real-time responses under human oversight).

Equally, some sovereign workloads remain best placed in national data centres, so long as they run on domestically governed infrastructure with audited pipelines.

Thales's edge-first architecture brings computing power and decision intelligence to the point of action, ensuring missions, services, and communications remain uninterrupted. Practically, our approach is policy-aware placement: deploy to the edge when mission continuity, bandwidth, or classification demands it; deploy to sovereign cloud/on-prem when centralisation improves assurance, scale or cost-efficiency, always under the same national governance and audit.

This approach is particularly critical in defence and aerospace contexts, where latency, cyber-resilience, and situational awareness are decisive.

3.3 SOVEREIGNTY THROUGH SECURE COLLABORATION

True sovereignty depends on a nation's ability to collaborate safely by sharing innovation and expertise without compromising on control over sensitive data, algorithms or infrastructure. In practice, this includes scenarios where nations may choose to share certain datasets or model components under bilateral or multilateral agreements, provided that such exchanges are governed, auditable and aligned with national policy.

Thales's approach to secure collaboration, refined through CortAIx and decades of R&D, is built on three complementary mechanisms: **transfer learning, collaborative learning, and synthetic data generation**. Together, these technologies make it possible to accelerate innovation collectively while preserving data sovereignty.

Transfer Learning: Sharing Models, Not Data

Transfer learning provides a balanced way to integrate global cooperation with national sovereignty. Shared foundational models, which are trained on public or anonymised datasets, can be adapted locally to national priorities using secure, sovereign data. While this approach reduces the need to exchange sensitive datasets, it still requires the right safeguards to prevent inversion risks.

To enhance interoperability and reduce vendor lock-in, Thales promotes open standards such as **ONNX (Open Neural Network Exchange)**, which allows models to be transferred seamlessly across frameworks and deployed within preferred infrastructures. By enabling

local control and transparency across the AI lifecycle, open standards strengthen sovereignty through technological independence.

Collaborative Learning: Federated and Split Approaches

Collaborative learning allows multiple organisations or agencies to train AI models together, without centralising their data. Each participant retains its own local dataset, exchanging only encrypted model updates rather than raw information whenever they do not want to share it.

The most mature form of this paradigm is **federated learning**, where all participants share a common model architecture that learns collectively across distributed environments. Emerging approaches such as **split learning** go a step further, dividing model components so that only intermediate activations, and not sensitive datasets, are exchanged.

Thales enhances these pipelines with cryptographic signing, provenance tracking and immutable model registries, ensuring that every contribution is authenticated, traceable and tamper-proof.

Synthetic Data: Privacy without Limiting Progress

Artificially generated yet realistic datasets allow training and validation where privacy or classification rules restrict the use of real-world data. These controlled synthetic environments make it possible to improve performance, diversity and robustness while maintaining strict compliance and security.

Through these secure collaboration frameworks, Thales enables nations to participate in shared innovation ecosystems, whether they choose to keep data fully local or share selected datasets under formal agreements with full assurance of sovereignty. Collaboration and control coexist by design, ensuring that progress is collective yet sovereign.

3.4 BUILDING BALANCE: THE DUAL APPROACH

For some nations, full sovereignty across the AI stack, from hardware to foundational models, remains a long-term challenge, particularly given global dependencies on advanced compute (such as GPUs).

Thales promotes a dual-sovereignty model:

- Develop proprietary, AI Sovereignty models for sensitive or strategic applications.
- Leverage open-source ecosystems for non-critical domains to maintain agility and innovation.

This approach enables countries to innovate rapidly while retaining strict control where it matters most, in critical sectors and decision chains. Thales promotes this model through modular architectures, where sovereign

and open components coexist safely within a trusted international framework.

4. Pillar Three – Innovation: Fueling National Capability and Growth

Innovation is the lifeblood of AI Sovereignty. It transforms ambition into capability, and capability into independence. For Thales, innovation is about empowering nations to develop, own and continuously evolve their own AI ecosystems.

By combining global expertise with local co-development, Thales helps countries transition from being technology adopters to becoming technology creators, driving sustainable growth through education, research, and partnership.

Thales invests in graduate pathways, upskilling programmes and joint research clusters to cultivate engineers, data scientists and policymakers who can design and assure AI. Co-development platforms localise capability, transfer know how and keep critical IP within national ecosystems, ensuring that innovation compounds domestically.

4.1 BUILDING THE FOUNDATIONS OF NATIONAL CAPABILITY

Sustainable innovation in AI Sovereignty depends on the co-evolution of four key enablers: data (curated, representative, and governed), infrastructure (compute, storage, and edge), talent (engineers, operators, and policymakers), and governance (ethics, safety, and compliance).

When these components move in harmony, nations build genuine digital independence; the ability to design, deploy and adapt AI on their own terms.

Among these, talent is the most critical. Compute and data can be procured; however, trusted expertise must be developed nationally. Thales recognises that human capital is a limiting factor in achieving sustainable AI sovereignty.

That is why Thales invests in graduate pathways, upskilling programmes and joint research clusters that cultivate the next generation of engineers, data scientists and policymakers. Through initiatives such as the Thales School of AI and CortAIx's specialised training modules, professionals across disciplines are equipped to design, deploy and govern AI responsibly.

In the UAE, Thales has collaborated with Sorbonne University Abu Dhabi for a hackathon focused on AI awareness and capability-building, an example of how academia-industry partnerships can inspire local innovation while reinforcing national strategy.

No single institution can deliver AI Sovereignty alone. True innovation happens when public, private and academic sectors work in concert, pooling expertise, aligning goals and ensuring that knowledge circulates within the national ecosystem.

Thales champions such public-private-academic partnerships, treating AI as a shared national asset rather than a proprietary technology. These collaborations accelerate research, validation, and deployment, ensuring that AI development serves both strategic objectives and societal needs.

Thales further operationalises the balance of openness and control through federated and transfer learning, which enable nations to collaborate securely without compromising data ownership. Models can learn from shared experience while data remains protected within national boundaries, proving that global innovation and national sovereignty can coexist.

For Thales in the UAE, collaborative innovation takes shape through TET as a platform for co-development, localisation and technology transfer. TET unites local engineering talent with Thales's global expertise to deliver innovation aligned with national priorities. It supports joint R&D, cultivates Emirati expertise, and ensures that critical know-how remains within the country.

This same model extends across Saudi Arabia, where Thales is seeding an ecosystem of localised engineering and academia-linked training, aligned with Saudi Vision 2030's focus on self-sufficiency and future-facing innovation. In both nations, the goal is to develop sovereign capability through people and partnerships.

4.3 TALENT: THE CORNERSTONE OF SUSTAINABLE INNOVATION

At the centre of every innovation strategy lies human potential. Thales's goal is not merely to deliver solutions but to empower people to develop and manage them independently. In the UAE, Thales aims to build teams composed of up to 30% national AI specialists, with similar efforts targeting 60% local talent in Saudi Arabia, creating a strong base of homegrown expertise. These internal goals reflect the deeper truth that sovereignty is sustained by national-level skill development.

Through the Thales Early Careers Programme and regional academic partnerships, the company is cultivating a new generation of AI professionals. They are equipped not only with technical proficiency but with an understanding of ethics, safety and governance. These individuals form the human core of AI Sovereignty, ensuring that capability grows from within, not through dependency.

4.2 COLLABORATION AS THE ENGINE OF INNOVATION

5. Pillar Four – Future Readiness: Anticipating the Future

If trust is the foundation of AI Sovereignty and innovation is its fuel, then future readiness is what ensures its longevity. Sovereignty must be a living capability, designed to evolve with every technological wave.

Modularity and scale allow nations to integrate new components without breaking certification or governance, and to move from pilots to national deployments under a single assurance framework.

Frugal AI delivers efficiency as a strength, and moving forward, post quantum cryptography and sovereign key management will be essential because quantum computing is already challenging today's encryption. Investing early in these capabilities ensures that AI architectures remain resilient long after today's technologies become legacy.

Thales defines future readiness as the ability to absorb, adapt and harness change before it dictates terms. It ensures that AI remains dependable, secure and aligned with human oversight, even as the environment shifts.

5.1 SOVEREIGNTY AS A LIVING CAPABILITY

AI sovereignty is a continuous state of preparedness. Technologies that were transformative a decade ago, from the rise of cloud computing to the current AI revolution, show how quickly national priorities can change. Future-ready nations must architect AI systems that are modular, upgradable and policy-responsive, allowing them to integrate new technologies without surrendering control or compliance.

Thales embeds this philosophy across its engineering approach through frameworks such as sovereign MLOps and AIOps toolchains, which are built for critical and even air-gapped environments.

Every artefact, whether data, code, or model, is digitally signed, version-controlled and auditable, ensuring complete traceability and trustworthiness. These policy-aware pipelines automatically enforce export-control rules, mission policies and security standards.

This approach treats sovereignty not as infrastructure, but as capability in motion that evolves as technology, regulation, and threats evolve.

5.2 ENGINEERING FOR MODULARITY AND SCALE

Future-ready AI must be both flexible and secure. Thales designs modular architectures that allow nations to integrate new components or algorithms without undermining certification or governance.

These architectures scale smoothly from local pilots to national-level deployments, ensuring that each addition, whether a new dataset, inference engine or analytics module, can be validated and governed within the same sovereign framework.

This modularity translates into hardware-agnostic performance. Thales's frugal, adaptive inference methods allow systems to maintain efficiency and reliability across diverse computing environments, from cloud to aircraft to battlefield. It sustains performance even under bandwidth limits, jamming, or energy constraints.

5.3 FRUGAL AI: EFFICIENCY AS STRENGTH

In critical systems, being smarter matters. Frugal AI is a defining differentiator for Thales, ensuring that AI performance remains robust even when data, compute, or power are limited.

Some airborne sensors, for instance, can generate over a terabyte of data per second, yet only a fraction of that can be stored or transmitted. Frugal AI focuses on selective retention and intelligent labelling, ensuring that only mission-critical insights are captured and processed.

Through techniques such as model quantisation, pruning, distillation, adaptive runtime and edge-first inference, Thales delivers AI that is fast, explainable, and resilient, capable of operating in real time with sub-100-millisecond responses under size, weight and power (SWaP) constraints.

This efficiency aligns with Net Zero objectives. By favouring localised computation and energy-efficient architectures, frugal AI significantly reduces the carbon footprint associated with massive, centralised training models. In Thales's design philosophy, sustainability and sovereignty reinforce each other, ensuring that energy-aware intelligence becomes a hallmark of responsible innovation.

5.4 PREPARING FOR THE NEXT TECHNOLOGICAL FRONTIER

Thales actively anticipates technologies that will redefine AI in the next decade. Among them, post-quantum cryptography (PQC) is a top priority. Quantum computing will eventually outpace today's encryption methods, demanding quantum-resistant algorithms and sovereign key management systems to secure future AI pipelines.

Another horizon technology is neuromorphic or brain-inspired computing, systems that mimic the neural efficiency of the human brain. These promise a breakthrough in on-edge learning and energy efficiency, dramatically reducing computational overhead while enabling real-time adaptation. Nations

that master both the hardware and algorithmic layers of this technology will hold a decisive advantage in performance, security, and sustainability.

By investing now in quantum-ready, hardware-agnostic and energy-adaptive systems, Thales helps partners ensure that their AI architectures will remain resilient long after today's technologies become legacy.

5.5 BALANCING OPEN INNOVATION WITH SOVEREIGN CONTROL

Future readiness also depends on the right balance between openness and control. Open-source technologies can strengthen sovereignty, but only when integrated within trusted, governed and dependency-vetted frameworks.

Thales's internal inner-source model follows this principle: an internal open-source ecosystem that accelerates secure collaboration and reuse across engineering teams while maintaining full auditability and oversight.

Externally, Thales's dual sovereignty approach combines agility with assurance. It enables nations to innovate rapidly while maintaining control where it matters most.

5.6 ETHICAL EVOLUTION AND HUMAN OVERSIGHT

As AI becomes more autonomous, ethics must evolve in parallel. Thales ensures that every system, regardless of domain or maturity, retains human oversight as a non-negotiable principle. Through its TrUE AI framework, the company embeds transparency, accountability and explainability into both engineering and operations.

Future-ready AI must be as responsible as it is resilient, capable of adapting to tomorrow's opportunities without compromising today's values. In doing so, AI Sovereignty becomes a societal asset: enhancing safety, strengthening accountability and building public trust while driving inclusive progress.

Ethical and societal readiness are integral to this vision. True sovereignty extends beyond technological control to reflect the values, culture and expectations of the societies it serves. Thales promotes governance models that align AI with national ethics and social frameworks, ensuring that trust is not only engineered, but also earned. In this way, AI Sovereignty reinforces the social contract between citizens, technology and the state.

services, Thales's solutions translate the four pillars of Trust, Sovereignty, Innovation, and Future Readiness into measurable gains in speed, precision, survivability, sustainability and workload reduction. The following examples demonstrate how AI can empower nations to act faster, decide smarter, and protect critical operations under full sovereign control.

GoldenAI (Electronic Warfare)

- **Outcome:** Up to 10× operator efficiency in post-mission analysis
- **Capability:** Automatically classifies 90% of intercepted radar emissions so operators can focus on the difficult 10%.
- **Sovereignty link:** Processes highly sensitive electronic-warfare signatures that must remain under national control.

An aircraft flying over a contested area intercepts multiple radar emissions. GoldenAI classifies the vast majority automatically, flagging only ambiguous signals for expert review. Analysts recover order-of-magnitude time savings, while raw EW data and signatures never leave sovereign infrastructure, preserving secrecy and independence.

Talios Pod (Airborne Target Recognition)

- **Outcome:** Up to 100× faster target detection and analysis; reduced crew workload and fatigue.
- **Capability:** Embedded AI analyses real-time imagery and automatically identifies targets, allowing pilots to focus on mission decisions.
- **Sovereignty link:** Deployed on national combat aircraft with AI qualified for defence-grade safety and security.

By accelerating image analysis in reconnaissance and combat missions, the Talios pod gives aircrew faster situational awareness and greater survivability while retaining human-in-the-loop control, illustrating trusted, frugal AI at work in critical systems.

6. Use Cases

Translating the Four Pillars into Real-World Impact

AI Sovereignty becomes clearest in the field. Across defence, cyber, maritime, aerospace, and public

Ground Master 200 Radar (Air Defence)



Outcome

3x improvement in UAV discrimination within cluttered airspace.

Capability

AI algorithms differentiate between low-speed targets (birds, drones, pedestrians) even in adverse weather.

Sovereignty link

Embedded AI trained and certified within national air-defence programs to ensure traceability and cybersecurity.



The outcome is faster, more reliable threat recognition for command centres, enhancing protection of national airspace and reducing operator cognitive load.

Ocean Control / Searchmaster Radar (Maritime Patrol)

- **Outcome:** 100x faster target recognition; improved situational awareness over wide areas.
- **Capability:** AI-based detection and classification of hundreds of maritime targets in seconds instead of minutes.
- **Sovereignty link:** Operated under national naval command structures to maintain information superiority.

Integrated aboard maritime patrol aircraft, this AI enhances efficiency and responsiveness, transforming data volume into decisive intelligence.

MiMap ATR (Maritime Warfare)

- **Outcome:** Automated target echo classification and contact-size estimation for faster mine countermeasure operations.
- **Capability:** AI assists operators by reducing manual interpretation time and increasing mission safety.
- **Sovereignty link:** Models trained on national acoustic datasets kept within secure naval infrastructure.

This capability improves under-sea situational awareness and supports faster decision-making in mission-critical scenarios.

Refit Optimizer (Maintenance and Fleet Planning)

- **Outcome:** Up to 10x faster maintenance planning; higher fleet availability and readiness.
- **Capability:** AI optimises scheduling and resource allocation across maintenance cycles.
- **Sovereignty link:** Deployed on national defence infrastructures to preserve data ownership and sustainment sovereignty.

By reducing administrative burden and turnaround time, Refit Optimizer enhances mission readiness and operational continuity.

GenAI4SOC (Cyber Operations)

- **Outcome:** Faster detection of complex cyber-attack patterns and improved analyst productivity.
- **Capability:** Generative AI models surface correlated anomalies and produce explainable incident reports in real time.
- **Sovereignty link:** Ensures critical security operations centres retain control over threat intelligence pipelines and response data.

This use case showcases how AI Sovereignty can fortify digital resilience without sacrificing transparency or human oversight.

TopSky Sequencer (Air Traffic Management)

- **Outcome:** 10x reduction in system configuration time; lower delays and fuel burn.
- **Capability:** AI optimises flight sequencing in real time, cutting months-long configuration down to days and reducing CO₂ emissions.
- **Sovereignty link:** Operated within national ATM systems with local data control and certification.

By combining operational efficiency with environmental responsibility, this application demonstrates how AI Sovereignty supports both economic and sustainability goals.

Frugal AI for Military Radio Communications

- **Outcome:** Improved audio clarity and 1000× lower power consumption than a standard laptop.
- **Capability:** AI-enhanced signal processing improves speech intelligibility in portable equipment.
- **Sovereignty link:** Co-developed with national defence agencies to ensure energy-efficient and secure communication under domestic control.

By embedding AI directly into soldier equipment, Thales reduces cognitive load and ensures continuous, resilient communications in the field.

Collaborative Combat Systems (Multi-Drone and Multi-Robot)

- **Outcome:** 10× increase in manageable fleet size and faster, more resilient operations.
- **Capability:** AI autonomously plans missions, prioritises tasks, and resolves trajectory conflicts for multi-agent systems.
- **Sovereignty link:** Algorithms designed for national defence networks to preserve control and compliance.

These systems reduce operator workload while boosting mission agility, embodying how trusted, AI Sovereignty augments human command without replacing it.

AI-Enhanced Activity-Based Intelligence (ABI) – MINDS

- **Outcome:** 4× faster intelligence analysis and 70% reduction in reporting time.
- **Capability:** AI-driven detection, recognition, and generative reporting transform raw imagery into actionable insights, revealing patterns of life and anomalies in strategic sites.
- **Sovereignty link:** Developed by Thales in compliance with NATO standards to ensure secure, autonomous intelligence processing for national defence forces.

By integrating advanced ABI and Generative AI into the MINDS imagery platform, Thales enables continuous, data-driven surveillance that strengthens situational awareness and decision-making superiority.

Thales's four pillars of Trust, Sovereignty, Innovation and Future Readiness provide a practical path to achieve that vision, powered by CortAlx, the Group's global AI accelerator that operationalises trusted, frugal and explainable AI for mission-critical use. Through CortAlx, Thales helps nations turn AI from imported dependency into a sovereign capability: secure, ethical and under national control.

The way forward is fundamentally clear: autonomy where it matters, collaboration where it benefits. AI Sovereignty is already delivering measurable impact with faster detection and analysis, lower cognitive load, greater survivability and reduced energy footprint, proving that defence-grade trust and edge-first design translate directly into mission advantage.

Sustainable sovereignty, however, begins with people. Engineers, operators and policymakers trained to design, assure and govern AI form the enduring human foundation of technological independence.

To accelerate progress from pilots to national platforms, governments and industry should:

- Establish national AI assurance frameworks aligned to standards such as ISO/IEC 42001, ensuring explainability, auditability and continuous testing.
- Operationalise end-to-end data sovereignty with jurisdictional control, access governance and federated training environments.
- Adopt dual sovereignty: sovereign models for sensitive workloads, vetted open-source stacks for non-critical domains.
- Prioritise frugal AI to guarantee continuity under constraint while cutting energy use.
- Invest in national talent pipelines through early-career programmes, retraining and specialist academies.

In this model, CortAlx stands as the bridge between global excellence and national capability, enabling nations achieve sovereignty through trust, transparency and performance at every layer of AI.

7. Conclusion

AI Sovereignty is a strategic necessity. Nations that can design, govern and evolve AI on their own terms gain decisive advantages in defence, resilience and competitiveness.



4, rue de la Verrerie 92190
Meudon FRANCE

Tél. + 33(0)1 57 77 80 00

www.thalesgroup.com

