POLICY & LEGAL FRAMEWORK · v1.0 · EFFECTIVE 31 MAY 2026

Research Policy

What we research — and what we don't.

Per-category legal framing for government, defence, medical, space, financial, climate, disaster, education, and agriculture research.

DCS AI Technologies L.L.C · License 1624450
Office 40, Dubai Industrial City, Saih Shuaib 3, Dubai, UAE · dcslabs.ai/research-policy
legal@dcsai.ai
research@dcsai.ai

Document control

Document	DCS Labs Research Policy
Version	v1.0
Effective date	31 May 2026
Owner	DCS AI Technologies L.L.C
Review cadence	Quarterly (next: 31 August 2026)
Document ID	DCS-LABS-POLICY-1.0
License	CC-BY-4.0 (this policy text); product code under MIT/Apache-2
Contact	legal@dcsai.ai · research@dcsai.ai

Contents

1. Why DCS Labs does this research

2. Per-category legal framing (9 sensitive domains)

01 Government · 02 Defence · 03 Medical · 04 Space · 05 Financial

06 Climate · 07 Disaster management · 08 Agriculture · 09 Education

3. Hard limits: things we will never build

4. Internal 5-step review process

5. Umbrella legal frameworks we comply with

6. Disclaimer of operational claims

7. How to challenge this policy

8. Glossary of legal references

1. Why DCS Labs does this research

DCS Labs is the research arm of DCS AI Technologies L.L.C, a UAE-licensed AI infrastructure company. We are not a government agency, not a defence contractor, not a clinical-grade medical-device manufacturer, and not a regulated financial institution. We are a sovereign-AI research lab whose 72-feature agenda spans the categories above. This document explains the legal frame we operate inside for each.

1.1 Sovereignty by default

If a customer's data leaves their jurisdiction, it loses much of its legal protection. Most AI infrastructure assumes US-controlled cloud. We research the opposite: cryptographically-verifiable sovereignty as a built-in property, not an add-on. This is the engineering and policy premise behind the R-Series (R+1 timestamp, R+2 provenance, R+3 audit export, R+4 zero-knowledge verification).

1.2 Auditability over hype

Every action, every receipt, every model decision, every agent move should be independently verifiable years later. Where the cryptography to do this honestly does not yet exist at production scale, we are honest about which receipts are still v0.1 research.

1.3 Open standards · closed weapons

Cryptography, receipt formats, agent protocols, and benchmark suites are published openly under MIT or Apache-2 (code) and CC-BY-4.0 (specs). Weaponizable applications are not built without sovereign partnerships, IRB-style ethical review, and explicit governmental authorisation. The two policies are complementary, not contradictory: open primitives, controlled deployments.

1.4 Research as a public good

Every paper, every prototype, every dataset, every benchmark we produce is published under permissive licenses. The 9-piece moat is in code; the standards are in specs; the 72-feature agenda is on this website. Build with us, fork us, audit us.

2. Per-category legal framing

Each of the nine sensitive categories below has its own legal frame. **Research mode** means whitepapers, prototypes, public demos, datasets, benchmarks, and reference code. **Operational mode** means deployment into the regulated environment, which requires partnership, certification, and authorisation we do not currently hold. Items currently in research-mode-only are marked clearly on the live 72-feature explorer at dcslabs.ai/research-72.

Domain 01 — Government

Scope: Research mode · **Sample features from the 72-feature roadmap:** #39, #47, #67 — sovereign AI stacks for ministries (India, GCC)

What we research

Sovereign-AI reference architectures for government use: air-gapped deployments, full audit chains, multi-ministry coordination patterns, citizen-data protection by cryptographic erasure. Published as whitepapers plus reference code that government IT teams may evaluate or fork. We accept government research-collaboration proposals where the deliverable is documentation or prototype, not operational systems.

What we explicitly do not do

- Operate government systems on behalf of any ministry.
- Bid on tenders for systems handling national-security data without a partnership-of-record and a defence-ministry-issued authorisation.
- Process classified information until and unless the relevant clearances are obtained for the specific deployment.
- Make claims of formal government partnership in marketing unless a publishable MoU exists.

Legal framework

UAE Federal Decree-Law 45/2021 (PDPL) · India DPDP Act 2023 · GDPR (EU residents) · Local government procurement law where engaged. Specific contracts may add classified-information handling clauses.

Domain 02 — Defence

Scope: Research mode · partnership-only · **Sample features from the 72-feature roadmap:** #8, #16, #20, #21, #22, #41, #46, #50 — dual-use AI primitives

What we research

Dual-use AI primitives that have a publishable, peer-reviewable, non-weaponized research surface: secure communication protocols, multi-agent coordination patterns, supply-chain optimisation models, threat-detection benchmarks (defensive). Research is published openly; deployments require a formal defence-ministry-authorised partnership-of-record.

What we explicitly do not do

- Autonomous weapon systems. Targeting components. Kill-chain integration. Autonomous strike-decision logic.
- Lethal-force-without-human-in-loop. Ever. There is no acceptable research framing for this.
- Deployment into operational defence networks without a sovereign defence partnership and explicit authorisation from the relevant defence ministry.
- Export of restricted-list technologies outside ITAR / EAR / Wassenaar permissions.
- Consulting on offensive cyber tooling for any actor, including state actors.

Legal framework

US ITAR (22 CFR Parts 120-130) · US EAR (15 CFR Parts 730-774) · Wassenaar Arrangement dual-use list · India SCOMET list (DGFT) · UAE export-control law (Federal Decree-Law 13/2007) · Local laws of any partner country.

> **IMPORTANT — defence research is partnership-gated.**
>
> If you are a defence prime, integrator, or government defence research organisation interested in any of the features in this category, the path is: (i) NDA + scope definition, (ii) classification of the work under your jurisdiction's export-control regime, (iii) authorisation letter from the appropriate ministry. We do not start defence work without these three steps.

Domain 03 — Medical / Healthcare

Scope: Research mode · not diagnostic · **Sample features from the 72-feature roadmap:** #25-27, #52-55, #72 — disease prediction, medical imaging, mental health

What we research

Privacy-preserving healthcare AI: disease-prediction models on anonymised data, medical-image analysis benchmarks, mental-health support agent architectures with sovereign data guarantees, drug-discovery acceleration prototypes. Output is papers, datasets, benchmarks, reference code. Pilots with hospitals or healthcare systems are research collaborations, not clinical deployments.

What we explicitly do not do

- Diagnostic output for clinical use. No 'this image shows cancer' classifications presented to a patient or clinician as actionable medical advice.

- FDA / CE / CDSCO-cleared product claims until the relevant clearance is obtained for the relevant product in the relevant jurisdiction.
- Mental-health crisis intervention as a primary care channel. Detection prototypes are research; deployments are referral-only, with a licensed professional human-in-loop, and never autonomous.
- Process Protected Health Information (PHI) in any form without a signed BAA (US), DPA (EU), or local equivalent.
- Pharmaceutical recommendations, dosage calculations, or interaction warnings presented as clinical advice.

Legal framework

HIPAA (US, with BAA) · FDA 21 CFR Part 820 + Software-as-Medical-Device (SaMD) guidance · EU Medical Device Regulation 2017/745 · India CDSCO Medical Devices Rules 2017 · GDPR Article 9 (sensitive data) · UAE Federal Law 2/2019 (health data).

Domain 04 — Space

Scope: Research mode · partnership-only · **Sample features from the 72-feature roadmap:** #7, #9, #14-15, #17-19, #23-24, #40, #42, #44-45, #48-49, #51, #68

What we research

Edge-AI architectures for spacecraft (low-power agents, intermittent connectivity), on-chain provenance for satellite imagery, mission-control assistant prototypes, multi-planet agent coordination protocols, lunar-resource-mapping AI, space-weather prediction networks. Open papers plus reference implementations. Engagements with space agencies are research-collaboration MoUs, not flight-software contracts.

What we explicitly do not do

- Operate space hardware. We are a software research lab; we do not fly anything.
- Claim ISRO, NASA, ESA, or other agency partnership unless a formal MoU is signed and publishable.
- Defense-adjacent satellite imagery analysis sold to private buyers without the relevant licensing.
- Flight-critical software for crewed missions without formal qualification regimes (DO-178C / NPR 7150.2 / equivalent).

Legal framework

UN Outer Space Treaty 1967 · US Remote Sensing Act + NOAA licensing · India Remote Sensing Data Policy 2011 · ITAR USML Category XV (spacecraft and satellites) · UAE Space Activities Law (Federal Law 12/2019).

Domain 05 — Financial

Scope: Research mode · not advice · **Sample features from the 72-feature roadmap:** #34, #35, #63, #64 — fraud detection, alternative credit scoring

What we research

Fraud-detection pattern research for banking and UPI, alternative credit-scoring on non-traditional data (with consent), insurance-fraud detection benchmarks. Models are published as research papers and open-source reference implementations; they are not deployed as personalised financial advice or as lending-decision systems in production without the relevant authorisations.

What we explicitly do not do

- Personalised financial advice. We are not a SEBI / RBI / FCA / SEC / DFSA-registered investment advisor.
- Payment processing as a Payment Institution. Our products use Stripe / SEPA / USDC; we are a SaaS, not a money-services business.
- Credit scoring as a service sold to lenders, until the relevant credit-reporting authorisations are obtained (RBI / equivalent) in the relevant jurisdiction.
- Use of synthetic identity for fraud-detection bypass research without explicit consortium-of-banks consent.
- Direct trading bots making market decisions with customer money outside a sandboxed test environment.

Legal framework

RBI Master Directions (India) · UPI Procedural Guidelines (NPCI) · PCI-DSS · EU PSD2 · UAE CBUAE regulations · SEC / SEBI / FCA / FINMA / DFSA depending on customer jurisdiction.

Domain 06 — Climate

Scope: Research mode · **Sample features from the 72-feature roadmap:** #30, #58, #59, #60 — climate risk, carbon credits, air quality, glaciers

What we research

Carbon-credit verification with satellite plus on-chain proofs, air-quality prediction and health alerts, glacier and sea-level monitoring, climate-risk prediction for farmers. Open-source benchmarks against public datasets (NASA EOSDIS, ESA Copernicus, India ISRO Bhuvan). We do not certify carbon credits ourselves; our research is verification infrastructure that recognised registries may adopt.

What we explicitly do not do

- Certify carbon credits as a registry. Our research is verification infrastructure that may inform Verra / Gold Standard methodologies.
- Sell catastrophe-bond pricing or insurance products without proper regulator engagement.
- Claim weather-modification or geoengineering capabilities.

Legal framework

UN Paris Agreement Article 6 · Verra VCS / Gold Standard methodologies (as references, not as a certifying body) · India BEE / MoEFCC notifications · EU Emissions Trading System (EU ETS).

Domain 07 — Disaster management

Scope: Research mode · NDMA-aligned · **Sample features from the 72-feature roadmap:** #31-33, #61, #62 — early warning, flood/cyclone, wildfire, earthquake, response coordination

What we research

Disaster early-warning models, flood and cyclone prediction, wildfire detection from satellite and drone imagery, earthquake-pattern research, emergency-response coordination protocols. Aligned with India National Disaster Management Authority (NDMA) and UN Sendai Framework reference frameworks. Pilots are research collaborations with agencies, not consumer-facing alerting services.

What we explicitly do not do

- Issue disaster warnings directly to populations. Warnings are issued by NDMA / IMD / national agencies; our research feeds them, not replaces them.
- Solo-actor emergency response automation. Humans-in-loop and licensed authorities always remain in the loop.
- Replace structural-engineering or seismological judgements with model output.

Legal framework

India Disaster Management Act 2005 · UN Sendai Framework 2015-2030 · WMO Common Alerting Protocol (CAP) · FEMA IPAWS (US, where applicable).

Domain 08 — Agriculture

Scope: Research mode · **Sample features from the 72-feature roadmap:** #28, #29, #56, #57 — crop monitoring, precision farming, soil, FCI supply chain

What we research

Crop-health monitoring from satellite and drone imagery, precision-farming agent recommendations, soil-health and fertiliser optimisation, supply-chain optimisation models for FCI-scale procurement. Open benchmarks; pilots with farmer cooperatives or government agricultural extension services where invited. We do not extract data from farmers without explicit, plain-language consent and a stated value-share.

What we explicitly do not do

- Agronomic prescriptions sold to farmers without local agriculture-extension partnership (KVK or equivalent).
- GMO or seed-engineering research — outside our remit.
- Data extraction from farmers without explicit, plain-language consent and a stated value-share.
- Pesticide or herbicide application recommendations without licensed-agronomist sign-off.

Legal framework

India Plant Quarantine Order 2003 · PPV&FR; Act 2001 (farmer rights) · EU CAP regulations · FAO Voluntary Guidelines on responsible tenure of land, fisheries, and forests.

Domain 09 — Education

Scope: Research mode · age-appropriate · **Sample features from the 72-feature roadmap:** #36, #37, #65, #66 — personalised learning, NEET/JEE, vernacular, skill gap

What we research

Personalised-learning agent architectures, skill-gap analysis models, NEET/JEE coaching agent prototypes, vernacular-language learning agents. All COPPA-aware; parental consent first for any user under 13. Deployments to school systems are pilot research, with teacher-in-loop decision-making.

What we explicitly do not do

- Data collection on minors without verified parental consent (COPPA-style framework applied even outside US jurisdictions).
- Autonomous grading or placement decisions presented as the school's official decision — teacher-in-loop always.
- Emotion analysis of students sold to administrators without published research showing harm-free use.
- Generate exam questions for exams the student has not yet sat (no cheating tools).

Legal framework

COPPA (US, under-13) · India DPDP Act 2023 (children's data) · GDPR Article 8 (consent age) · UNESCO AI in Education recommendations · Local education-ministry rules where deployed.

3. Hard limits — things we will never build

Some research areas are technically possible but ethically and legally off-limits. These will never enter our roadmap. If a customer, partner, or government agency asks for them, the answer is no, even if commercial terms are offered. This list is not exhaustive but represents our standing red lines.

01 · Autonomous weapon systems

Targeting, kill-chain integration, autonomous strike decisions, lethal-force-without-human-in-loop. Not researched, not prototyped, not consulted on. Includes defensive autonomous systems whose failure mode is lethal.

02 · Mass surveillance toolchains

Face-recognition plus cross-referencing-without-warrant pipelines. Population-scale tracking. Social-credit scoring. Universal biometric databases for state actors.

03 · Biometric inference at scale

Gender, sexuality, political-affiliation, religious-belief, or health-status inference from face, voice, or other biometric signals. No exceptions, including "academic" framings.

04 · Undisclosed AI in deceptive contexts

Generative voice or text representing a real person without their explicit, recorded consent. Synthetic media in elections. Romance-scam AI. Impersonation AI in legal, medical, or banking contexts.

05 · Bio / chem / nuclear weapon precursor research

Any AI assistance to biological, chemical, or nuclear weapons design, synthesis pathways, delivery mechanism, or evasion of monitoring systems. Includes academic-curiosity prompts.

06 · Child-safety violations

CSAM generation. Detection-bypass research. Age-verification-bypass research. We report any such request to the relevant authorities and to the NCMEC CyberTipline (US) or equivalent.

07 · Critical-infrastructure attack research

Offensive cyber-attack tooling targeting power-grid, water, hospital, transport, financial, or telecommunications infrastructure. We may research defence; never offence.

08 · Election-influence operations

Voter-targeting automation, disinformation generation, deep-fake political ads, persuasion-as-a-service targeting electoral outcomes. Includes paid consultancy on this topic.

4. Internal 5-step review process

Every feature on the 72-item roadmap that touches one of the nine sensitive categories above goes through this 5-step internal review before code is written. Items matching the hard-limits list in Section 3 are auto-rejected at step 1 and recorded in our refusal log.

Step 1 · Category classification

Each proposed feature is classified by domain (Government / Defence / Medical / Space / Financial / Climate / Disaster / Agriculture / Education / Core Infrastructure). The hard-limits list is checked first — any match auto-rejects with a brief written reason that is logged.

Step 2 · Legal-framework mapping

The relevant statutes are listed (e.g. HIPAA + FDA SaMD for medical, ITAR + EAR for defence). If we lack the relevant authorisation, the feature is downgraded to "research-mode only" — no operational deployment, even on request from a customer or partner.

Step 3 · Harm model

What is the worst-case misuse? Who is the affected population? What harm-mitigation primitives must be in the design (consent, opt-out, audit, kill-switch, rate-limit)? If we can't name a meaningful mitigation, the feature is downgraded to "research-mode only" pending design iteration.

Step 4 · Disclosure plan

How will we be honest about what is built versus theoretical? What status pill (Shipped / Partial / Roadmap) is accurate? What flags ship default-OFF? We do not advertise capabilities that are not operationally enabled.

Step 5 · Reversibility

If after launch we discover unintended harm, can we roll back? Cryptographic erasure, kill-switches, model unloads, and consent revocation must work — proven before launch, not after. If reversibility is not possible, the feature does not ship.

5. Umbrella legal frameworks

Beyond category-specific law, DCS Labs operates under these umbrella frameworks for every product, every customer, every jurisdiction.

Data protection

GDPR (EU + UK), India DPDP 2023, UAE Federal Decree-Law 45/2021, US state laws (CCPA / CPRA / VCDPA), Brazil LGPD. DPA template at dcsai.ai/dpa. Sub-processor list at dcsai.ai/sub-processors.

Export control

US ITAR + EAR, Wassenaar Arrangement, India SCOMET (DGFT), UAE export-control law. No dual-use technology export without licence verification. End-use and end-user screening for any restricted-list technology engagement.

Intellectual property

Our code is published under MIT or Apache-2 licences. Our standards (R-Series, Trust SKU) under CC-BY-4.0. We respect third-party IP; takedown process at legal@dcsai.ai. No use of unlicensed training data.

AI-specific regulation

EU AI Act risk-tier compliance (high-risk = research-mode only until certification). NIST AI Risk Management Framework as reference. India NITI Aayog Responsible AI principles. UAE AI Charter. Voluntary G7 Hiroshima principles.

Security

Responsible-disclosure process at dcslabs.ai/security. Bug-bounty programme for receipt-chain and R-Series vulnerabilities. PGP key at dcslabs.ai/security/pgp.txt. Quarterly third-party security review.

Anti-corruption

US FCPA, UK Bribery Act 2010, India Prevention of Corruption Act, UAE Federal Decree-Law 31/2021. Zero tolerance for facilitation payments. Annual employee training. Confidential ethics line at ethics@dcsai.ai.

Trade and sanctions

OFAC SDN list screening for all customers. UK HMT Consolidated List. EU sanctions list. UAE targeted financial sanctions list. No service to designated entities.

Tax and accounting

IFRS reporting. UAE Corporate Tax Law (Federal Decree-Law 47/2022). VAT registered. Statutory audit annual.

6. Disclaimer of operational claims

This policy is normative for what we research, not declarative about what is deployed. Many items on the 72-feature roadmap are research-mode-only and will remain so unless and until operational authorisations are obtained for the relevant deployment context.

Nothing in this policy, on the dcslabs.ai website, or in our whitepapers should be interpreted as:

- An offer of regulated services in any jurisdiction where DCS AI Technologies L.L.C is not authorised.
- Medical, legal, financial, tax, or accounting advice.
- A claim of operational defence, space, or government capability beyond what is explicitly stated, with the appropriate authorisation, in a signed agreement.
- A guarantee that any research-mode feature will become a product on any specific timeline.
- A commitment to enter any specific market or category beyond the published roadmap.

Investors, partners, regulators, and customers should rely on signed agreements and published audit-grade receipts (R+2 / R+3) for assurances. Marketing claims, blog posts, social-media activity, and aspirational roadmap statements are not contractual.

7. How to challenge this policy

If you believe a feature on our roadmap should not exist, tell us. If you believe we are being too cautious in a category that matters to you, also tell us. If you find an inconsistency between this policy and our public website, code, or product behaviour, especially tell us. We update this policy quarterly and immediately whenever a partner raises a question we had not previously answered.

Channels

- **General policy questions:** legal@dcsai.ai
- **Research partnership inquiries:** research@dcsai.ai
- **Security disclosure:** security@dcsai.ai or dcslabs.ai/security
- **Ethics concerns:** ethics@dcsai.ai (confidential)
- **Press inquiries:** press@dcsai.ai
- **Postal address:** Office 40, Dubai Industrial City, Saih Shuaib 3, Dubai, UAE

Response commitments

- Acknowledgement within 3 working days.
- Substantive response within 10 working days for non-trivial questions.
- Policy update commitment within 30 working days if the question reveals a gap.
- Annual transparency report summarising all material questions received and our response.

8. Glossary of legal references

Short-form references used in this document. Full citations available on request.

Term	Definition
AI Act (EU)	Regulation (EU) 2024/1689 on harmonised rules on artificial intelligence.
BAA	Business Associate Agreement (HIPAA-required for PHI handling).
CCPA / CPRA	California Consumer Privacy Act / California Privacy Rights Act.
CDSCO	Central Drugs Standard Control Organization (India medical devices).
COPPA	Children's Online Privacy Protection Act (US, under-13).
DGFT	Directorate General of Foreign Trade (India export licensing).
DPA	Data Processing Agreement (GDPR-required contract).
DPDP Act	Digital Personal Data Protection Act 2023 (India).
EAR	Export Administration Regulations (US dual-use export control, 15 CFR 730-774).
FCPA	Foreign Corrupt Practices Act (US anti-bribery).
FDA SaMD	US Food and Drug Administration Software-as-Medical-Device guidance.
GDPR	General Data Protection Regulation (EU 2016/679 + UK retained).
HIPAA	Health Insurance Portability and Accountability Act (US).
ITAR	International Traffic in Arms Regulations (US, 22 CFR 120-130).
LGPD	Lei Geral de Proteção de Dados (Brazil data protection).
MeitY	Ministry of Electronics and Information Technology (India).
NDMA	National Disaster Management Authority (India).
OFAC SDN	US Office of Foreign Assets Control Specially Designated Nationals list.
PDPL (UAE)	Personal Data Protection Law (UAE Federal Decree-Law 45/2021).
PSD2	Revised Payment Services Directive (EU).
RBI	Reserve Bank of India (banking regulator).
SCOMET	Special Chemicals, Organisms, Materials, Equipment and Technologies (India export-control list).
Wassenaar	Wassenaar Arrangement on Export Controls for Conventional Arms and Dual-Use Goods.