

Ziru Labs Capability Posture

A Reference Primer

Ziru Labs Research Publication · Version 1.0 · 2026 · Published under Ziru Labs corporate byline with contribution from Daniel Martin · Distributed under Creative Commons Attribution 4.0 International (CC BY 4.0)

Citable reference: Ziru Labs. *Ziru Labs Capability Posture: A Reference Primer*, v1.0. Published at zirulabs.com/research.

ABSTRACT

This primer documents Ziru Labs' capability posture within the trust layer for AI category at category level. It describes the threat classes the company's substrate addresses (mapped against the Physics-Layer Threat Taxonomy for AI Infrastructure), the composition patterns through which the substrate connects to adjacent categories (per the Trust Layer Category Map), the deployment surface posture across regulated commercial and federal classified contexts, the compliance and cleared-handling posture under which the company operates, and the intellectual property, engineering, partnership, and standards-engagement postures that situate the company within the category.

The primer is the operational complement to the analytical foundation papers (The Trust Layer for AI: A Reference Framework for the Category; The AI Infrastructure Stack and the Trust Layer Position: A Reference Primer; The Physics-Layer Threat Taxonomy for AI Infrastructure; The Runtime Verification Gap in Federal AI Deployment: A Reference Primer; and The Trust Layer Category Map: A Reference Primer; all Ziru Labs, 2026). The analytical foundation establishes the category and its landscape. This primer documents Ziru Labs' specific posture within both.

The document operates under deliberate export-control discipline. It describes what Ziru Labs addresses and what it composes with at category level. It does not describe how the substrate operates at the mechanism level. This boundary is structural rather than presentational: mechanism-level detail at the trust-layer substrate sits within export-controlled technology classifications, and the company's standing discipline is to make the operational posture legible to counterparties without crossing the export-control boundary.

The primer is structured to support due diligence by counterparties evaluating engagement with Ziru Labs, by readers evaluating the category's commercial reality against its analytical claims, and by standards-body and sovereign-program participants situating Ziru Labs within their respective frameworks.

1. Purpose of This Document

The analytical foundation papers establish the trust layer for AI category, its position in the AI infrastructure stack, the threat landscape it addresses, the operator landscape adjacent to it, and the federal deployment context in which it currently matters most. The Operator Profile sections of those papers identify Ziru Labs as the publisher of the analytical foundation and as one operator within the category, but those sections are brief by editorial design: the analytical papers are deliberately neutral reference documents that should stand on their own merits independent of the company that authored them.

This primer is the place where Ziru Labs documents its specific posture within the category. It is intended to answer the questions that a sophisticated counterparty conducting due diligence on Ziru Labs would ask after reading the analytical foundation: what does the company address against the threat taxonomy it published, how does the company compose with the adjacent operators its category map surveyed, what is the deployment surface posture across the contexts the federal deployment primer described, what is the compliance and intellectual property posture, and what is the engagement architecture for counterparties.

The primer is not a marketing document. The brand discipline that governs the analytical papers also governs this primer: McKinsey register, neutral institutional voice, no claims that exceed what the corpus can substantiate at category level. Where specific implementation detail would be useful to a reader but sits within export-control restriction, the primer says so explicitly rather than substituting marketing language for substance.

2. The Disclosure Boundary

The trust-layer substrate operates at the silicon-and-chassis layer of AI infrastructure. Mechanism-level detail at this layer sits within multiple overlapping export-control classifications, including the U.S. Export Administration Regulations sections governing dual-use cryptographic and computing technology, the U.S. International Traffic in Arms Regulations sections governing defense articles where applicable, and the Wassenaar Arrangement multilateral export control framework. The company operates under the assumption that mechanism-level disclosure could constitute a controlled disclosure regardless of the disclosure venue; the standing discipline is therefore to operate at the category level in all public-facing materials.

The practical implication is that this primer describes the substrate in terms of what it addresses (threat classes, application surfaces, deployment contexts), what it composes with (adjacent operators and standards), what compliance posture it operates under, and what posture the company maintains across intellectual property, engineering, and partnerships. The primer does not describe how the substrate works at silicon level, what specific mechanisms operate at the physical layer, what specific signal pathways are used, what specific cryptographic constructions are employed, or what specific accelerator-vendor or platform-vendor implementations are in flight.

This discipline is read by some readers as overcaution and by others as appropriate posture. Both readings exist. The company's position is that the export-control posture is not negotiable in public-facing materials, and that the analytical foundation papers plus this primer together provide sufficient substance for any counterparty to evaluate engagement without the company needing to cross the export-control boundary publicly. Counterparties requiring mechanism-level evaluation operate under the appropriate cleared-handling arrangements, which the Engagement section addresses.

Readers coming to this primer from analytical work in other deep-tech categories may find the discipline more restrictive than is conventional. The cleanest framing is that the trust layer for AI category, by virtue of operating at the silicon-and-chassis layer with explicit defense, intelligence, and sovereign-AI deployment relevance, sits structurally closer to the export-controlled end of the deep-tech spectrum than do AI software, AI applications, or even most AI infrastructure categories above silicon. The discipline reflects that structural reality.

3. Threat Class Coverage Against the Taxonomy

Ziru Labs' substrate is designed to address, in whole or in part, threat classes across all five surfaces. The coverage posture below is described at category level using the taxonomy's threat class identifiers; the specific mechanism by which the substrate addresses each threat class is not described. Coverage is stated at the level of architectural design: the substrate is pre-commercial, with a minimum working prototype targeted for the second half of 2026, and the coverage posture describes what the architecture is designed to do rather than capability demonstrated in a shipped product.

Attack Surface 1: Physical Access. Coverage across all four threat classes documented in the taxonomy. S1.T1 Cold-Boot Memory Extraction is addressed by design through hardware-rooted memory protection that operates against post-power-state memory contents. S1.T2 Chip-Level Physical Probing is addressed by design through hardware-rooted assumptions about what silicon contents remain extractable outside the operational envelope. S1.T3 Chassis-Level Tampering is addressed by design through hardware-rooted tamper response with zeroization triggers at the chassis layer. S1.T4 Supply Chain Compromise Post-Manufacture is addressed by design through hardware-rooted component-level attestation supporting supply-chain integrity verification.

Attack Surface 2: Bus and Interconnect. Coverage across the four threat classes documented in the taxonomy. S2.T1 PCIe Bus Scraping Under Operating System Compromise, S2.T2 NVLink and High-Bandwidth Fabric Interception, S2.T3 CXL and Memory-Pool Fabric Interception, and S2.T4 Lateral Memory Traversal Across Tenant Boundaries are each addressed by design through hardware-rooted bus-level mechanisms that operate at the layer below the operating system's ability to observe or control. The specific mechanism by which the substrate addresses these threats varies by interconnect type and is not described in this primer.

Attack Surface 3: Firmware and Boot. Partial coverage. S3.T2 Driver-Level Privilege Escalation Affecting AI Inference Integrity, S3.T3 Boot-Chain Compromise, and S3.T4 Runtime Firmware Modification are addressed by design through hardware-rooted attestation that does not depend on driver or firmware integrity. S3.T1 GPU Firmware Injection or Modification falls within the silicon-vendor confidential computing domain and is addressed through composition with the silicon-vendor implementations surveyed in the Category Map; Ziru Labs does not duplicate the silicon-vendor confidential computing coverage of this threat class.

Attack Surface 4: Inference Integrity. Coverage across all four threat classes documented in the taxonomy. S4.T1 Hardware-Level Inference Manipulation is addressed by design through hardware-rooted fault-detection and inference-integrity verification. S4.T2 Inference Result Substitution is addressed by design through hardware-rooted cryptographic signing of inference outputs. S4.T3 Software-Layer Compliance Bypass is addressed by design through hardware-rooted compliance attestation that operates below the software-layer policy enforcement mechanisms. S4.T4 Model Weight Extraction via Inference Side Channels is addressed by design through hardware-rooted weight protection and inference observation isolation.

Attack Surface 5: AI Governance Persistence. Coverage across all three threat classes documented in the taxonomy. S5.T1 Jailbreak Persistence Across Operating System Compromise is addressed by design through hardware-rooted policy enforcement with session-independent state. S5.T2 Alignment Constraint Software Bypass is addressed by design through hardware-enforced constraint enforcement operating below the alignment-

mechanism software layer. S5.T3 Governance Enforcement Decoupling is addressed by design through hardware-rooted attestation that operates independently of the governance mechanism being attested.

The coverage posture above is the company's current published position. The Threat Taxonomy is a living reference document, and the substrate's coverage will evolve with the taxonomy through subsequent versions of both documents.

What the substrate does not address. Two categories of threat that confront AI infrastructure are explicitly out of scope for the Ziru Labs substrate at the time of this writing. First, adversarial machine learning at the model layer (prompt injection prevention at content-classifier level, adversarial input detection, membership inference defense, and adjacent model-layer concerns) is addressed by other operators in the category landscape (AI security and governance platforms at Layer 11, per the Category Map). The substrate composes with those operators rather than substituting for them. Second, conventional kernel and hypervisor security is addressed by the established information security industry and is out of scope for the substrate, which assumes those layers as adjacent prerequisite content per the Threat Taxonomy's scoping section.

4. Composition with Adjacent Categories

The substrate composes additively with operators at adjacent layer positions as documented in The Trust Layer Category Map (Ziru Labs, 2026). The composition postures below describe how Ziru Labs' substrate connects to each adjacent category. The composition is structural; specific composition with named operators within each category proceeds through commercial and standards-body arrangements rather than through public posture.

Composition with hardware roots of trust at Layer 4. The substrate composes with the established hardware root of trust ecosystem (Trusted Platform Module ecosystem, Apple Secure Enclave, Google Titan family, Microsoft Pluton, AWS Nitro Security Chip, ARM TrustZone, OpenTitan) by extending the attestation envelope from boot-time and configuration-time measurement to continuous runtime evidence of AI workload operation. The hardware root of trust establishes the integrity baseline; the substrate produces runtime evidence against that baseline.

Composition with confidential computing at Layer 5. The substrate composes with the silicon-vendor confidential computing implementations (NVIDIA Confidential Computing, Intel TDX, AMD SEV-SNP, ARM CCA) and with the pure-play confidential computing operators (Anjuna, Fortanix, Edgeless Systems, and others) by extending the attested surface below the trusted execution environment boundary. Confidential computing protects the confidentiality of workload data within an active session, including in-die memory and encrypted bus traffic, and produces runtime attestation that the platform is genuine and the environment unmodified; the substrate provides hardware-rooted evidence and enforcement of the AI computation itself, across the physical-layer, bus-level, and supply-chain surfaces and the session-boundary windows that confidential-computing attestation does not reach. The two compositions together produce defense-in-depth trust posture from outside-the-chassis through inside-the-die: confidential computing establishes that the environment is genuine, and the substrate establishes what the computation did and under which constraints.

Composition with AI security and governance platforms at Layer 11. The substrate composes with the AI security and governance platform ecosystem (HiddenLayer, Robust Intelligence, Lakera, Protect AI, CalypsoAI, Credo AI, Holistic AI, Fiddler, Cranium, IBM watsonx.governance, and others) by providing hardware-rooted enforcement and evidence at the substrate level that the platforms cannot produce from their own software layer. Governance platforms provide the breadth of monitoring and policy-management coverage across enterprise AI estates; the substrate provides the depth of hardware-rooted enforcement at the inference verification surface. The two compositions together produce governance posture that survives software-layer compromise.

Composition with hardware security specialists. The substrate composes with the hardware security IP ecosystem (Rambus, Synopsys DesignWare, Cadence), hardware Trojan and supply-chain integrity specialists (Cycuity and adjacent operators), anti-tamper component vendors (Analog Devices, Microchip Technology), and security-oriented open-architecture silicon vendors (SiFive, OpenTitan). The composition with these specialists proceeds at the silicon-design and silicon-supply-chain layers and is part of the substrate's deployment lifecycle architecture.

Composition with cryptographic verifiable inference. The substrate composes with the cryptographic verifiable-inference operators surveyed in the Category Map (the zero-knowledge machine learning and proof-of-inference ecosystem, including EZKL, Lagrange Labs, Inference Labs, and Nexus) by rooting their software-generated proofs in attested silicon. Cryptographic verifiable inference proves, in software, that a committed model produced a given output; the substrate establishes that the computation ran on attested hardware within an attested physical envelope, and can bind the software proof to that hardware-rooted evidence. The substrate's own hardware-rooted signing of inference outputs (Attack Surface 4, threat class S4.T2) and the cryptographic proof systems are complementary rather than substitutive: the combination produces a verifiable-inference claim that is both mathematically sound and physically rooted, across the physical-layer and bus-level surfaces the cryptographic abstraction does not reach.

Composition with standards bodies and sovereign programs. The substrate composes with the standards codification activity at ISO/IEC JTC 1/SC 27 (ISO/IEC 27090), ISO/IEC JTC 1/SC 42, CEN-CENELEC JTC 21, NIST (AI RMF, Cybersecurity Framework, FIPS), the Confidential Computing Consortium, and the Trusted Computing Group. Ziru Labs participates in the codification activity through the standards-body engagement posture described in Section 8. The substrate composes with the sovereign-program landscape (DARPA, IARPA, DoD CDAO, allied program offices) through the partnership architecture described in Section 8 and within the cleared-handling posture described in Section 6.

The composition posture across all six is consistent with the Reference Framework's positioning of the trust layer as a substrate that operates by enabling rather than displacing adjacent categories.

5. Deployment Surface Posture

The substrate is designed to deploy across the full range of AI infrastructure contexts where hardware-rooted runtime verification is required. The deployment surface posture below is described at category level; specific deployment configurations are subject to counterparty-specific scoping under the engagement architecture in Section 10.

Federal classified deployment. The substrate is designed to deploy in U.S. federal classified AI infrastructure at Impact Level 6 and above per the Defense Information Systems Agency Cloud Computing Security Requirements Guide. The Runtime Verification Gap in Federal AI Deployment: A Reference Primer (Ziru Labs, 2026) documents the deployment context. The substrate is designed to compose with the FY2026 NDAA Section 1513 (Physical and Cybersecurity Procurement Requirements for Artificial Intelligence Systems, P.L. 119-60) framework as that framework codifies through Department of Defense implementation activity.

Allied sovereign deployment. The substrate is designed to deploy in allied sovereign AI infrastructure at classifications equivalent to U.S. Impact Level 5 and above, including United Kingdom Government Security Classifications Policy SECRET and above, NATO classification tiers from NATO SECRET upward, and equivalent allied classifications across Five Eyes, NATO member states, and Gulf Cooperation Council member states. The substrate is designed to compose with NATO STANAG AI trust frameworks as those frameworks codify and with allied national frameworks at structurally equivalent positions.

EU regulated commercial deployment. The substrate is designed to deploy in EU regulated commercial AI infrastructure under EU AI Act high-risk and successor application categorizations. The substrate is designed to compose with conformity assessment under EU AI Act Articles 40 (harmonised standards) and 43 (conformity assessment procedures) as the harmonised standards under CEN-CENELEC JTC 21 codify.

Frontier laboratory deployment. The substrate is designed to deploy in frontier AI laboratory infrastructure where responsible scaling policies require hardware-rooted demonstration of safety constraints during inference and training. Specific frontier laboratory engagement proceeds under the partnership architecture described in Section 8 and is subject to specific commercial arrangement.

Regulated vertical commercial deployment. The substrate is designed to deploy in regulated vertical commercial AI infrastructure including financial services AI under fiduciary requirements, healthcare AI under HIPAA and regulatory analogs, legal services AI under burden-of-proof requirements, and adjacent regulated verticals where hardware-rooted evidence of AI computation is required for audit or compliance demonstration.

The deployment surface posture is broader than the substrate's launch configuration. Launch deployment proceeds in the federal classified and allied sovereign contexts first, with subsequent extension into the EU regulated commercial, frontier laboratory, and regulated vertical contexts as commercial configurations mature.

Application surface extension. The substrate's deployment posture extends across the application surfaces identified in the companion Reference Framework (Ziru Labs, 2026): security, verification, provenance, compliance, identity, and economic agency. Security application maturity is driven by the U.S. federal classified-deployment frameworks documented in the Runtime Verification Gap primer (Ziru Labs, 2026), by FY2026 NDAA Section 1513, by ISO/IEC 27090, and by allied sovereign frameworks. The verification, provenance, compliance, identity, and agency application surfaces are driven by the U.S. Center for AI Standards and Innovation (CAISI) at NIST, the Coalition for Content Provenance and Authenticity (C2PA), EU AI Act Article 50 machine-readable disclosure requirements taking effect August 2, 2026, FedRAMP 20x continuous monitoring rollout, and adjacent sector-specific frameworks including ABA Formal Opinion 512 for legal services, FINRA Regulation BI for financial services, FDA AI/ML SaMD guidance and HIPAA per-claim PHI handling for healthcare, and federal records management requirements. The substrate is designed to compose with the

technical requirements emerging from each of these codification activities; specific application surface mechanisms are documented under cleared engagement protocols and are not publicly disclosed during the active standards-codification window. Across the provenance and agency surfaces in particular, the substrate's contribution is the hardware-rooted binding beneath software-layer agent-action attestation and content-provenance assertions: it roots such claims in attested silicon rather than generating them at the software layer, and composes with the software-layer operators surveyed in the Trust Layer Category Map (Ziru Labs, 2026) rather than substituting for them.

6. Compliance and Cleared-Handling Posture

The substrate operates within compliance and cleared-handling postures appropriate to the deployment contexts identified above.

U.S. federal classified handling. The company operates under the cleared-handling discipline appropriate to engagement with U.S. federal classified AI infrastructure. Specific cleared-handling arrangements are subject to standard federal cleared-handling protocols and are not disclosed publicly. Counterparties operating under matching cleared-handling protocols engage through the appropriate channels.

Export-control posture. The company operates under U.S. Export Administration Regulations compliance discipline appropriate to the technology classifications under which the substrate sits. The category-level disclosure discipline applied across the company's public-facing materials is one operational expression of the export-control posture. Counterparties evaluating engagement with the substrate in contexts requiring export-control evaluation engage through the appropriate channels under Section 10.

EU AI Act conformity posture. The substrate is designed to operate under EU AI Act conformity assessment regimes as the harmonised standards codify. The company participates in CEN-CENELEC JTC 21 working group activity as appropriate to the substrate's positioning within the harmonised standards work.

Standards alignment. The substrate is designed to align with FIPS 140-3 physical security requirements (Levels 3 and 4 as deployment context requires), NIST SP 800-193 platform firmware resiliency requirements, NIST AI Risk Management Framework GOVERN and MAP function requirements, NIST Cybersecurity Framework 2.0 functional requirements, ISO/IEC 42001:2023 AI management system requirements, ISO/IEC 23894:2023 AI risk management requirements, and the in-progress ISO/IEC 27090 AI security guidance as that standard codifies. Specific alignment claims at deployment level are subject to deployment-specific evaluation and certification under the appropriate frameworks.

Cloud and platform certification posture. The substrate is designed to operate within FedRAMP authorization boundaries at Moderate and High impact levels, within DISA Impact Level authorization boundaries at IL5 and IL6, and within equivalent allied authorization frameworks at structurally equivalent positions. Specific deployment certifications proceed through deployment-specific evaluation under the relevant authorization framework.

The compliance posture is structured to allow the substrate to deploy across regulated, classified, and sovereign contexts without requiring per-context substrate redesign. The architectural posture supporting this is described at

category level; specific compliance attainment is deployment-specific.

7. Intellectual Property Posture

The substrate's intellectual property position is structured around mechanism-level patent claims developed across the relevant operating layers of the substrate. The patent position is in active development across multiple filing jurisdictions appropriate to the substrate's deployment contexts. Specific patent claim detail is not disclosed publicly during the active filing period; the company's standing position is that disclosure of patent claim detail outside the patent filing process is not part of the public-facing posture.

The intellectual property posture also includes trade secret protection across silicon-design, mechanism-level, and operational know-how that is appropriately maintained as trade secret rather than as patent claim. Trade-secret-protected material is by definition not disclosed publicly; the existence of the trade secret discipline is a feature of the company's IP architecture rather than a description of specific material.

The intellectual property pipeline is actively expanding across three primitive-extension vectors: application-domain extension into adjacent domains where the substrate's mechanism category produces value beyond the security launch deployment; category-defense extension preserving the substrate's position as the category-defining mechanism holder as the physics-layer threat landscape and the broader AI compute stack evolve; and capability extension through complementary hardware and cryptographic mechanisms that unlock additional substrate deployments. The forward pipeline is structured to expand as the substrate's applications are mapped into domains beyond the launch deployment, consistent with the substrate-category development pattern documented in the Reference Framework. Specific invention detail across the forward pipeline is not disclosed publicly during the active filing period; the existence and structural direction of the pipeline is documented here to make the company's substrate-extension posture legible at category level.

The combination of patent claims at mechanism level and trade secret protection across operational know-how is the standard intellectual property architecture for deep-tech substrate-category operators at the silicon-and-chassis layer. Historical analogs in the Reference Framework (ARM, Qualcomm, VeriSign in their respective category-formation periods) each employed comparable IP architectures during their substrate-formation phases.

The IP posture is one of the structural reasons the Layer 3 position in the AI infrastructure stack is sparsely populated at the time of this writing, per the Category Map's observation. Substrate-category positions at this layer require mechanism-level patent development that does not emerge through software-pattern capital and engineering structures.

8. Engineering and Development Posture

The substrate is in active engineering development. The development posture is described at category level; specific engineering milestones, headcount, and burn are not part of the public-facing posture in alignment with the standard deep-tech development discipline at this category formation stage.

Engineering composition. The engineering posture includes silicon-design competence, hardware-software integration competence, cryptographic protocol design competence, and the cleared-handling capacity required for engagement with federal classified and allied sovereign deployment contexts. Specific engineering team composition is not disclosed publicly.

Development partnerships. The substrate is designed to compose with silicon-vendor platforms and with confidential computing implementations as documented in Section 4. Specific silicon-vendor and platform-vendor development partnerships proceed through commercial arrangements that are subject to mutual disclosure terms. Public partnership disclosure proceeds at the time of mutual announcement.

External research review. Ziru Labs' technical design documentation is under review by a university-affiliated government-security research program, under non-disclosure agreement. The engagement is at the documentation-review stage; no findings, validation, or endorsement are claimed or implied.

Standards-body engagement. The company is pursuing standards-body activity relevant to the substrate's deployment contexts, including engagement with ISO/IEC JTC 1/SC 42, ISO/IEC JTC 1/SC 27, CEN-CENELEC JTC 21, IEEE working groups, the Confidential Computing Consortium, the Trusted Computing Group, the U.S. Center for AI Standards and Innovation (CAISI) at NIST, the Coalition for Content Provenance and Authenticity (C2PA), and adjacent bodies at the appropriate participation levels. The standards-body engagement portfolio spans hardware-rooted infrastructure trust (Confidential Computing Consortium, Trusted Computing Group, ISO/IEC JTC 1/SC 27), AI management and risk frameworks (ISO/IEC JTC 1/SC 42, NIST AI RMF profile work), regional regulatory implementation (CEN-CENELEC JTC 21 for EU AI Act harmonised standards under Articles 40 and 43), and emerging agent-interoperability and reasoning-provenance frameworks (NIST CAISI Agent Interoperability Profile, C2PA v2.4 assertion types), where the company's engagement is at the hardware-rooting layer beneath software-layer agent-attestation and provenance methods rather than at the attestation layer itself. Future standards-body submissions operationalize the analytical framework developed in the corpus and are in active preparation across multiple bodies.

Sovereign program engagement. The company anticipates engagement with sovereign-program activity relevant to the substrate's deployment contexts, including engagement with U.S. federal program offices and with allied program offices at appropriate participation levels. Specific sovereign-program engagements proceed under the cleared-handling postures appropriate to each program and are not disclosed publicly except where mutual disclosure has been agreed.

Founding team. The substrate's development draws on the founding team's prior experience spanning U.S. federal classified AI, signals intelligence, and cryptographic warfare, including senior service in the U.S. intelligence community and the U.S. Navy, alongside experience in financial services and AI infrastructure. The founding team's composition is appropriate to the substrate's deployment contexts. Future versions of the corpus will carry named contributions from members of the team and from external contributors as the company's public posture develops; specific founding team disclosure proceeds through the company's information on zirulabs.com and through specific counterparty engagement under the architecture described in Section 10.

9. Partnership and Standards Architecture

The substrate's commercial deployment proceeds through a partnership architecture composed of three primary partnership categories.

Silicon-vendor partnerships. The substrate composes with the silicon-vendor compute and confidential-computing layer as documented in Section 4. Silicon-vendor partnerships proceed under standard commercial arrangements appropriate to substrate-category integration, including joint development, reference architecture co-publication, and deployment co-engagement where contexts warrant. Specific silicon-vendor partnership disclosure proceeds at mutual announcement.

Platform-vendor and cloud partnerships. The substrate is designed to deploy on cloud and platform infrastructure including the major cloud providers' federal and commercial offerings. Platform-vendor and cloud partnerships proceed under standard commercial arrangements appropriate to substrate deployment, including integration engineering, certification co-engagement, and joint customer engagement where appropriate. Specific partnership disclosure proceeds at mutual announcement.

Sovereign and federal-systems-integrator partnerships. The substrate's deployment in U.S. federal classified and allied sovereign contexts proceeds through partnership with established federal-systems-integrator counterparties and through direct sovereign-program engagement where appropriate. Specific partnership detail is subject to the cleared-handling postures appropriate to each engagement.

The partnership architecture is structured to allow the substrate to deploy at the scale and across the contexts the analytical foundation papers identify, without requiring the company to substitute for the established silicon-vendor, platform-vendor, and federal-systems-integrator capabilities that the substrate composes with.

10. What This Capability Posture Is Not

The posture document is explicit about what it does not provide.

It is not a product specification. Specific deployment specifications proceed through counterparty-specific scoping under the engagement architecture below. The posture document describes coverage and composition at category level; product-level specification is deployment-context-specific.

It is not a financial or commercial disclosure. Specific commercial terms, pricing, financing, valuation, headcount, burn, and adjacent commercial metrics are not part of the public-facing posture. Counterparties evaluating engagement under commercial arrangements engage through the architecture below.

It is not a roadmap document. Specific engineering and product roadmap detail is not part of the public-facing posture in alignment with standard deep-tech development discipline at this category formation stage. The corpus identifies the application surfaces the substrate is designed to serve (security, verification, provenance, compliance, identity, agency) and the order in which deployment proceeds (security first, others as commercial configurations mature). Specific roadmap milestones beyond this category-level posture are not publicly disclosed.

It is not a competitive document. The Trust Layer Category Map (Ziru Labs, 2026) surveys the operator landscape at category level. This posture document describes Ziru Labs at category level within that landscape. Neither document makes competitive claims about Ziru Labs versus other operators; the analytical posture across the corpus is structural rather than competitive.

It is not a substitute for cleared-handling engagement. Counterparties evaluating engagement with the substrate in contexts requiring cleared-handling, export-controlled, or otherwise restricted disclosure engage under the appropriate protocols. The posture document is the public-facing artifact; substantive engagement under restricted-handling protocols proceeds through the architecture below.

11. How to Engage

Counterparties evaluating engagement with Ziru Labs route through the Engage pathway at zirulabs.com. The Engage pathway accommodates three categories of counterparty engagement.

Commercial counterparties evaluating engagement under standard commercial arrangements (silicon-vendor partnerships, platform-vendor partnerships, federal-systems-integrator partnerships, regulated commercial deployment) engage through the appropriate Engage channel with scoping conversations proceeding under standard mutual non-disclosure terms.

Cleared counterparties evaluating engagement under cleared-handling protocols (federal classified, allied sovereign, intelligence community, defense industrial base) engage through the appropriate Engage channel with scoping conversations proceeding under the cleared-handling protocols mutually appropriate to the engagement.

Standards-body and academic counterparties evaluating engagement under standards-body or academic protocols (joint standards-body work, joint academic research, joint publication) engage through the appropriate Engage channel with scoping proceeding under the relevant standards-body or academic protocols.

General research inquiries route to research@zirulabs.com. Standards-body submissions and academic correspondence route through the appropriate Engage channel.

The Engage pathway is the canonical entry point for counterparty engagement. The pathway is structured to route counterparties to the appropriate engagement protocol without requiring the counterparty to navigate the company's internal organization.

12. Bibliography

References are organized into five categories: academic literature (A), standards and regulatory documents (B), industry and vendor technical documentation (C), historical and analytical reporting (D), and framework integration references (E).

A. ACADEMIC LITERATURE

The analytical foundation for this posture document draws on the academic literature catalogued in the bibliographies of the companion analytical foundation papers. Direct academic citation is concentrated in those papers; this posture document

references the analytical foundation rather than duplicating the academic citation.

B. STANDARDS AND REGULATORY DOCUMENTS

European Commission (2024). “Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act).” Official Journal of the European Union.

Executive Order 14179 (2025). “Removing Barriers to American Leadership in Artificial Intelligence.” The White House, signed 23 January 2025; 90 FR 8741.

FY2026 National Defense Authorization Act, Section 1513 (2025). “Physical and Cybersecurity Procurement Requirements for Artificial Intelligence Systems.” Public Law 119-60.

International Organization for Standardization / International Electrotechnical Commission (2023). “ISO/IEC 42001:2023, Information technology, Artificial intelligence, Management system.”

International Organization for Standardization / International Electrotechnical Commission (2023). “ISO/IEC 23894:2023, Information technology, Artificial intelligence, Guidance on risk management.”

International Organization for Standardization / International Electrotechnical Commission (2026, Final Draft International Standard). “ISO/IEC 27090, Cybersecurity, Artificial intelligence, Guidance for addressing security threats to artificial intelligence systems.”

National Institute of Standards and Technology (2018). “NIST SP 800-193: Platform Firmware Resiliency Guidelines.”

National Institute of Standards and Technology (2019). “FIPS 140-3: Security Requirements for Cryptographic Modules.”

National Institute of Standards and Technology (2023). “AI Risk Management Framework (AI RMF 1.0),” NIST AI 100-1.

National Institute of Standards and Technology (2024). “AI RMF Playbook.”

National Institute of Standards and Technology (2024). “The NIST Cybersecurity Framework (CSF) 2.0,” NIST CSWP 29.

U.S. Export Administration Regulations, 15 CFR Parts 730-774 (current).

Wassenaar Arrangement on Export Controls for Conventional Arms and Dual-Use Goods and Technologies (current).

C. INDUSTRY AND VENDOR TECHNICAL DOCUMENTATION

Confidential Computing Consortium (current). Technical documentation and member working group output.

Trusted Computing Group (current). “TPM 2.0 Library Specification” and “DICE Architecture.”

Detailed industry and vendor technical documentation for the operators with which the substrate composes is catalogued in the bibliography of *The Trust Layer Category Map: A Reference Primer* (Ziru Labs, 2026).

D. HISTORICAL AND ANALYTICAL REPORTING

Ziru Labs (2026). *The Trust Layer for AI: A Reference Framework for the Category, v1.0*. Published at zirulabs.com/research.

Ziru Labs (2026). *The AI Infrastructure Stack and the Trust Layer Position: A Reference Primer, v1.0*. Published at zirulabs.com/research.

Ziru Labs (2026). *The Physics-Layer Threat Taxonomy for AI Infrastructure: A Reference Framework, v1.0*. Published at zirulabs.com/research.

Ziru Labs (2026). *The Runtime Verification Gap in Federal AI Deployment: A Reference Primer, v1.0*. Published at zirulabs.com/research.

Ziru Labs (2026). *The Trust Layer Category Map: A Reference Primer, v1.0*. Published at zirulabs.com/research.

E. FRAMEWORK INTEGRATION REFERENCES

Defense Information Systems Agency (current revision). “Cloud Computing Security Requirements Guide.”

MITRE Corporation (2024). “ATT&CK Framework, Enterprise Matrix,” Version 14.

MITRE Corporation (current). “ATLAS: Adversarial Threat Landscape for Artificial-Intelligence Systems.”

North Atlantic Treaty Organization (current). “NATO Standardization Agreement (STANAG) framework activities on AI trust.”

UK National Cyber Security Centre (current). “Government Security Classifications Policy.”

Acknowledgments

This primer draws on the analytical foundation established in the companion research publications, the standards-and-regulatory work currently active across the trust layer for AI category, and the founding team’s prior experience spanning U.S. federal classified AI, signals intelligence, cryptographic warfare, financial services, and AI infrastructure. Specific acknowledgments to the Ziru Labs founding team. Any errors in characterization are solely Ziru Labs’ responsibility.

Citation

Ziru Labs. *Ziru Labs Capability Posture: A Reference Primer, v1.0*. Published at zirulabs.com/research.