Strategic Business Plan and Pitch Architecture: FABIABox and the Sovereign AI Paradigm

Executive Summary:

FABIABox & The SquadShelf Ecosystem

Company: SquadShelf (Spin-off from IABAI Research) **Founder:** Andrea Sannuto **Sector:** Sovereign AI Infrastructure / Hardware **Target Raise:** $1.5M - $2.5M Pre-Seed (SAFE)

The Vision: "Company as a System"

We are moving the corporate world beyond Software-as-a-Service (SaaS). FABIABox is the industrial-grade hardware backbone for the Agentic Economy. It is the "3D Printer for Corporations," allowing visionaries to seamlessly print, deploy, and manage autonomous, specialized AI agent squads locally on sovereign silicon.

The Problem: The Sovereign AI Bottleneck

The enterprise AI revolution is stalling at the execution layer. Cloud-based AI models represent a massive data-sovereignty liability and are fundamentally incompatible with the strict data localization mandates of the **EU AI Act of 2026**. Furthermore, consumer AI wrappers lack the reliability and legal standing to execute real corporate workflows. To run autonomously, businesses need absolute data privacy, local processing, and recognized legal agency.

The Solution: FABIABox & Foundry OS

FABIABox is a plug-and-play desktop server powered by cutting-edge silicon (NVIDIA Thor / AMD Ryzen AI Max+). It runs a hardened, immutable NixOS-based kernel—**Foundry OS**—ensuring zero data leakage to public clouds.

FABIABox is not just hardware; it is the execution node for the complete **SquadShelf Ecosystem**, integrating 27 years of IT framework R&D:

- **Intelligence Gathering:** *IntentRanker* and *IdeaTally* map and prioritize corporate goals.
- **Data Refinery:** *DocuOstream* captures unstructured data, and *IAB.AI* uses it to locally fine-tune specialized models (e.g., Gemma 4).
- **Legal Agency:** *SquadSign* utilizes our proprietary "Legal Slicing" protocol and Restricted Power of Attorney, allowing agents to execute transactions and sign contracts legally under **eIDAS 2.0** guidelines.

Market Opportunity

Enterprise buyers are aggressively seeking Compliant-by-Design, air-gapped infrastructure. The global Sovereign AI segment is valued at **$8B to $12B**, expanding rapidly due to European regulatory pressures. By combining localized compute with digital legal agency, FABIABox captures an immediate **$2.5B Total Addressable Market (TAM)** at the exact inflection point of enterprise AI adoption.

Business Model: Dual-Engine Revenue

We monetize the entire autonomous workflow, blending high-margin physical sales with recurring digital micropayments:

1. **Hardware (CapEx):** Direct unit sales of the FABIABox (est. $5,000/unit), establishing our physical footprint inside the enterprise.
2. **Utility (OpEx):** Recurring micro-payments via Model Context Protocol (MCP) services (like LayerExtract and DocuOstream) utilized by agents.
3. **Licensing:** Subscription fees for SquadSign's legal API and eIDAS compliance node.

The Ask: $1.5M - $2.5M Pre-Seed SAFE

The foundational software is built, the legal rails are set, and the IABAI spin-off is underway. We are opening a Pre-Seed round to accelerate global hardware tooling, fulfill initial production runs, and scale go-to-market operations.

FABIABox offers early-stage investors a rare opportunity to own the critical, legally compliant infrastructure layer of the $50B autonomous workforce market.

Macro-Economic Market Analysis: The Sovereign AI Surge (2026)

The trajectory of the global technology market indicates a decisive, capital-intensive pivot

away from centralized data centers toward edge computing and localized AI inference. This shift is driven by the prohibitive bandwidth costs of transmitting massive volumes of unstructured data to the cloud, the microsecond latency requirements of real-time agentic workflows, and the escalating demand for data sovereignty imposed by national governments.

Sizing the Global AI and Edge Computing Infrastructure Market

Current market valuations project exponential, sustained growth within the edge computing sector. Conservative estimates indicate that the global edge computing market reached $33.4 billion in 2025 and is projected to scale to $328.0 billion by 2033, representing a Compound Annual Growth Rate (CAGR) of 32.1%.[6] More aggressive financial models forecast the market scaling from $658.1 billion in 2026 to $1.86 trillion by 2031, fueled primarily by the aggressive deployment of edge infrastructure to support low-latency data processing and localized AI-enabled workloads.[7]

The broader global AI infrastructure market already exceeds $50 billion annually, compounding at a 30% rate as enterprise procurement shifts from software licenses to underlying compute capacity.[1] Within this expansive context, the specific segment for Sovereign AI—defined strictly by hardware and software ecosystems that prioritize localized data control, intellectual property protection, and regulatory compliance—is actively expanding. The addressable segment for localized, edge-based AI currently sits at an $8 billion to $12 billion global valuation.[1]

Industry analysts universally note that 2026 marks the definitive "year of sovereign AI." Significant infrastructure investments are accelerating across heavily regulated industries such as healthcare, financial services, defense, and legal services.[2] Enterprises in these sectors are moving their mission-critical AI workloads into sovereign environments to retain absolute control over models and governance.[2]

The FABIABox targets the highly lucrative intersection of developer-first AI tools, pre-seed capital allocators, and enterprise digital transformation officers. This specific cross-section represents a validated $2.5 billion TAM.[1] The initial Serviceable Obtainable Market (SOM) for the FABIABox's hybrid hardware-software model is concentrated in the developer and accredited investor segments, accounting for $40 million to $60 million in annualized contract value (ACV) and equity deployment capacity.[1]

Market Segment Focus	2025/2026 Estimated Valuation	2030/2033 Projected Valuation	Projected CAGR	Primary Growth Drivers
Global Edge Computing	$33.4B (2025) / $46.7B (2026)	$328.0B (2033)	32.1% [6]	Latency reduction, localized processing, network

				resilience.
Global AI Infrastructure	>$50.0B (2025)	N/A	30.0% [1]	Model training, automated enterprise agentic workflows.
Edge AI Software & Hardware	$37.5B (2026)	$102.9B (2030)	28.7% [8]	Real-time processing, privacy mandates, autonomous deployment.
Sovereign AI Sub-Segment	$8.0B - $12.0B (2026)	N/A	N/A [1]	Data localization mandates, EU AI Act compliance, IP protection.
FABIABox Intersection TAM	$2.5B (2026)	N/A	N/A [1]	Plug-and-play desktop nodes, early-stage AI investment flows.

The Fracturing of the "Software Seat" Licensing Paradigm

For two decades, the SaaS industry has relied heavily on the "Software Seat" licensing model, wherein enterprises pay monthly subscriptions based on the number of human users interacting with the software. Agentic AI inherently dismantles this financial paradigm. When an autonomous agent operates 24/7 without human intervention, traditional per-seat licensing fails to capture the value generated and simultaneously introduces prohibitive, highly unpredictable API token costs.

Furthermore, cloud-based agentic workflows exhibit severe structural vulnerabilities that enterprise chief information security officers (CISOs) are actively rejecting. When a human utilizes a SaaS tool, the human acts as the final arbiter of data verification and contextual appropriateness. However, when an autonomous agent is granted system access, it operates

independently; a hallucination or an unverified tool call executed in the cloud can result in catastrophic data leakage, compliance breaches, or systemic database corruption.[5]
Enterprise buyers are highly cognizant of this "Centralized Cloud Hallucination & Leakage Risk." Consequently, the deployment of agentic automation—which market intelligence firm IDC projects will benefit over 40% of enterprise applications by 2026—requires an infrastructural foundation built on absolute cryptographic predictability.[2] The transition from software-as-a-service to "Company as a System" necessitates that the hardware compute, the operating system, the language model, and the agentic orchestrator function as a single, air-gapped, and immutable unit residing on the enterprise's physical premises.

The Critical Bottleneck: Regulatory Compliance and Data Sovereignty

The timing of the FABIABox product launch coincides perfectly with a historic and uncompromising tightening of digital regulations across the European Union. Products engineered to be "Compliant-by-Design" for the European market possess a distinct and highly lucrative competitive advantage. This compliance acts as an impenetrable economic moat against foreign cloud competitors that rely on fluid data borders and opaque data supply chains.

The EU AI Act of 2026: The Imminent Enforcement Catalyst

The EU Artificial Intelligence Act establishes the world's most comprehensive and punitive regulatory framework for artificial intelligence systems. While recent legislative amendments have introduced staggered enforcement delays for specific high-risk product categories—pushing certain Annex III systemic obligations (e.g., AI used in critical infrastructure, law enforcement, and employment) to December 2027—the core transparency obligations and foundational compliance requirements become fully enforceable on August 2, 2026.[3]
For developers and enterprises targeting the European market, treating these regulatory delays as permission to pause compliance planning is universally recognized as a severe strategic error.[12] The technical file required under Article 11 of the EU AI Act is not a simple administrative checklist; it requires a robust body of verifiable documentation demonstrating the system's design decisions, training data governance, and fundamental rights impact assessments.[12] Existing approaches to regulatory compliance remain largely manual and fragmented, making them highly vulnerable to audit failures.[13]
The FABIABox architecture anticipates and systematically neutralizes these compliance hurdles by executing entirely on local silicon. Because no unstructured corporate data is transmitted to third-party cloud inference providers, the enterprise retains absolute end-to-end control over the data supply chain. This fundamentally alters the risk calculus under the EU AI Act. The system structurally prevents "single-token leaks," ensuring that sensitive corporate data utilized for local agentic fine-tuning remains strictly within the organization's legal perimeter, thereby satisfying the most rigorous interpretations of data sovereignty.

eIDAS 2.0 and the Digital Legal Agency of Agents

A critical historical bottleneck in the deployment of autonomous agents is the fundamental lack of recognized legal agency. An AI language model can draft a highly complex legal contract, but it cannot traditionally execute it, sign it, or bind an organization to its terms. This severe limitation is resolved natively through the intersection of the FABIABox software stack and the newly implemented eIDAS 2.0 regulations.

The European Digital Identity Act (eIDAS 2.0, formally Regulation (EU) 2024/1183) mandates that all member states offer highly secure digital identity wallets to their 450 million citizens by the end of 2026, and legally requires that private-sector enterprises, banks, and healthcare providers accept these wallets by late 2027.[4] The framework introduces Qualified Electronic Attestations of Attributes (QEAAs), which are verifiable digital statements issued by trusted organizations.[4] Crucially, under eIDAS 2.0, a Qualified Electronic Signature (QES) carries the exact legal equivalence of a handwritten signature and mandates cross-border recognition across all EU member states.[16]

The FABIABox operationalizes this regulatory breakthrough through a proprietary legal slicing protocol named "SquadSign." SquadSign acts as a localized digital notary mechanism deeply embedded within the hardware's execution loop. By utilizing eIDAS 2.0 compliant cryptographic signatures and verifiable credentials, an enterprise executive can issue a "Restricted Power of Attorney" (POA) token directly to an autonomous agent housed within the FABIABox.[17]

This architectural innovation transforms the agent from a passive text generator into an active, legally recognized entity capable of securely executing financial transactions, verifying Know Your Customer (KYC) identity data, and signing commercial contracts.[18] This native integration of digital legal rails constitutes a formidable, insurmountable barrier to entry for cloud-native competitors, who cannot guarantee the cryptographic isolation and localized security required for qualified digital signatures.

EU Regulatory Framework	Implementation Deadline	Core Mandate	Impact on Cloud AI	FABIABox Strategic Solution
EU AI Act (Transparency)	August 2, 2026 [3]	Disclosure of AI generation and user interactions.	High liability for opaque third-party API wrappers.	Complete local execution ensures total algorithmic transparency.
EU AI Act (High-Risk	December 2, 2027 [10]	Rigorous data governance, human	Cloud data transmission severely	Air-gapped "Compliant-by-Design"

Annex III)		oversight, and risk management.	complicates compliance audits.	architecture neutralizes data transmission risks.
eIDAS 2.0 (Digital Wallets)	End of 2026 [4]	Mandatory provision of EUDI wallets to citizens.	Cloud agents lack the local secure enclaves for identity custody.	Native integration with secure local hardware for credential storage.
eIDAS 2.0 (QES & QEAA)	December 2027 (Acceptance)	Cross-border legal validity of digital signatures and attributes. [4]	Cannot securely hold Restricted Power of Attorney.	SquadSign protocol enables legal slicing and binding agent transactions.

Technical Architecture: The Foundry of Sovereign Silicon

The physical hardware and underlying software operating system of the FABIABox represent the true industrialization of local AI. To execute complex, highly parameterized foundation models alongside an overarching autonomous orchestrator without reliance on cloud compute, the device leverages state-of-the-art unified memory architectures and immutable operating environments. The hardware tier relies on high-margin unit sales, presenting a compelling Capital Expenditure (CapEx) alternative to perpetually escalating cloud Operating Expenditure (OpEx).

Comparative Silicon Economics: AMD and NVIDIA Paradigms

The current desktop AI workstation market is defined by an intense architectural battle between NVIDIA's established market dominance and AMD's highly disruptive unified memory implementations.

NVIDIA's strategic approach is exemplified by the DGX Spark, a compact desktop AI supercomputer powered by the Grace Blackwell GB10 superchip.[20] It fuses an ARM-based Grace CPU featuring up to 20 cores with a Blackwell GPU featuring 6,144 CUDA cores and 128GB of LPDDR5X unified memory, offering approximately 1 PetaFLOP of parallel throughput.[20] Built explicitly for local AI inference, agent prototyping, and data science workflows, the DGX Spark operates on a custom Ubuntu Linux stack and is priced at

approximately $4,699, having seen recent price escalations due to global supply constraints affecting LPDDR5X memory.[20]
Conversely, AMD's newly released Strix Halo architecture, branded as the Ryzen AI Max+ 395, offers a highly disruptive, cost-efficient alternative.[25] This platform features 16 desktop-class Zen 5 CPU cores, a dedicated XDNA 2 NPU capable of 50 TOPS, and a truly massive integrated GPU driven by 40 AMD RDNA 3.5 Compute Units (Radeon 8060S configuration).[25] Crucially, it supports up to 128GB of LPDDR5X unified memory operating on a giant 256-bit bus, of which up to 96GB can be dynamically converted to Video RAM (VRAM) through AMD Variable Graphics Memory.[25] Systems utilizing this architecture are entering the market at approximately $3,999, undercutting NVIDIA's offering by roughly $700 while offering native support for both Windows 11 and Linux environments.[23]
The FABIABox specifically leverages these advanced unified memory topologies to achieve the local execution of massive foundation models. The 128GB of memory is the foundational element that allows the system to overcome the traditional VRAM bottleneck that previously restricted complex local AI to expensive, multi-GPU server racks.[28]

The Data Loop: Local Fine-Tuning and Model Viability

The availability of 96GB to 128GB of usable VRAM allows the FABIABox to natively run state-of-the-art open-weight models, transforming unstructured corporate data into precise execution. The proprietary IAB.AI pipeline is engineered to utilize these models dynamically. While Google's Gemma 3 family is heavily optimized for these environments, possessing a dense transformer architecture and massive 128K token context windows [29], the FABIABox infrastructure extends capabilities further. Utilizing the IAB.AI pipeline, the development team has successfully fine-tuned highly advanced architectures, including targeted iterations of the Gemma 4 framework, directly on highly complex, unstructured datasets locally.
This proves the core architectural thesis: the FABIABox can build hyper-specialized expert agents entirely on-device. For context regarding hardware demands, advanced models approaching the 27-billion parameter scale require approximately 54GB of VRAM in unquantized BF16 format, and roughly 15.1GB when utilizing highly efficient Q4_K_M quantization.[29] With a 128GB unified memory pool, the FABIABox can comfortably host these massive models alongside localized orchestrators, multimodal vision encoders, and localized vector databases without paging to slower storage drives. Furthermore, utilizing highly optimized tools like Unsloth, the hardware allows enterprises to conduct local fine-tuning up to 16x faster while maintaining 128k token context lengths.[31] This capability validates the FABIABox's primary moat: turning proprietary data into specialized execution without external network exposure.

Software Architecture: Immutable Orchestration and OpenShell Sandboxing

Raw compute is merely the prerequisite; the software layer is what fundamentally transforms the hardware into an enterprise-grade digital factory. Autonomous agents possess the inherent

capability to execute shell code, access file systems, and communicate over internal networks. In a traditional operating system architecture, this creates an unacceptably vast security attack surface.[9] The FABIABox neutralizes this systemic risk through the deployment of an immutable kernel and rigorous containerized sandboxing.

The underlying operating system of the FABIABox relies on a hardened, NixOS-based kernel. NixOS represents a radical departure from traditional Linux distributions; it relies entirely on declarative configuration and immutable infrastructure.[33] Every aspect of the system—from critical packages and background services to networking rules and user permissions—is expressed strictly as code and activated atomically.[34]

For an edge appliance operating autonomously like the FABIABox, immutability is paramount. The core OS files are mounted as strictly read-only within the nix/store, establishing an environment where unauthorized modifications, whether initiated by external malicious actors or rogue AI agents exhibiting hallucinated behaviors, are systematically prevented at the kernel level.[35] If a system update is required, the NixOS architecture builds a completely new configuration in an isolated path before seamlessly switching over. This wholly eliminates configuration drift across decentralized enterprise deployments and mathematically guarantees that the system always boots into a known-good, verifiable state.[33] This level of operational predictability and auditability is highly attractive to enterprise compliance officers navigating the stringent documentation requirements of the EU AI Act.[12]

To orchestrate the AI agents securely atop this immutable foundation, the software stack heavily integrates NVIDIA NemoClaw (and the underlying open-source OpenClaw framework). NemoClaw functions as a comprehensive, open-source reference stack specifically designed for deploying domain-specialized, always-on AI agents that reason, plan, and act.[37]

A critical, non-negotiable component of this orchestration stack is the OpenShell sandboxing environment. OpenShell operates as an advanced policy engine that enforces uncompromising privacy and security guardrails around the executing agent.[5] Rather than granting an agent unfettered access to system resources or internal APIs, OpenShell utilizes low-level Linux security technologies—specifically Landlock, Seccomp, and strict Network Namespaces—to definitively gate agent behaviors.[9]

Within the FABIABox, the core business orchestrator—dubbed FABIA—utilizes OpenShell to command specialized sub-agents. If a sub-agent attempts to transmit sensitive financial data or intellectual property to an unauthorized external cloud endpoint, the OpenShell network policy intercepts and blocks the request at the kernel level.[5] This cryptographic guarantee ensures that sensitive data stays entirely within the company's network, satisfying the deepest legal requirements of data sovereignty and providing absolute assurance to European corporate buyers seeking to avoid compliance liabilities.

Go-To-Market Strategy and the Dual-Engine Revenue Model

The commercial monetization strategy for the FABIABox deliberately transitions value capture away from highly elastic, static software subscriptions toward a highly lucrative, dual-engine

framework. This model flawlessly synergizes high-margin capital expenditure (the Hardware Tier) with a continuous, infinitely scalable transaction-based operating expenditure (the Modular Utility Tier).

Tier 1: CapEx Hardware Economics and Trust Signaling

The physical FABIABox unit represents the entry point to the ecosystem. Positioned competitively at an MSRP of $5,000 per unit, the hardware is sold as a high-margin enterprise capital expenditure [user prompt]. This pricing intentionally undercuts traditional centralized server deployments while providing equivalent, localized inferential power.

Crucially, selling the device as a physical hardware appliance bypasses the inherent skepticism currently associated with static software delivery.[1] It leverages powerful, tangible trust signals; enterprise procurement departments and CISOs are historically conditioned to procure physical hardware for security, compliance, and localized operations. This makes the procurement cycle significantly smoother for risk-averse departments evaluating AI tools against GDPR, the EU AI Act, and emerging data sovereignty frameworks.[1]

Tier 2: OpEx, Model Context Protocol (MCP), and Micropayments

Once the hardware is successfully deployed, the ecosystem generates massive, recurring operational revenue through the Modular Utility Tier. This tier is built entirely upon the Model Context Protocol (MCP). MCP is an emerging open standard championed by foundational model providers that creates a unified interface for AI models to securely call external tools, fetch contextual data, and interact with services.[39] It functions effectively as the "API revolution" for agentic AI.[39]

The FABIABox ecosystem operates a standalone suite of agent-agnostic MCP services, deploying proprietary tools such as "LayerExtract" for highly accurate unstructured data ingestion, and "DocuOstream" for interfacing with the SquadSign legal notary protocol. These services act as digital tollbooths within the localized and wider network. While the core MCP standard has seen rapid global adoption, the ecosystem is currently plagued by a lack of commercial reliability; many free MCP servers lack SLA guarantees, enterprise security compliance, or dedicated support, creating a massive void for premium, highly reliable paid alternatives.[40]

FABIABox aggressively capitalizes on this void. Any internal enterprise agent, or explicitly authorized third-party AI system operating within the network, can access these specialized MCP tools on a per-call basis.

To facilitate the monetization of this MCP utility tier seamlessly, the ecosystem integrates advanced machine-to-machine (M2M) micropayment architectures, heavily inspired by the mechanics of networks like TollBit and M2M Escrow smart contracts.[42] In a paradigm where AI agents execute thousands of micro-tasks daily, standard $20/month software subscriptions are highly inefficient and mathematically flawed.[43]

Instead, the infrastructure demands a mechanism where an agent can securely pay fractions of a cent per API call.[44] Platforms like TollBit enable this dynamic by deploying specialized MCP endpoints that meter usage, validate calling agents via JSON Web Tokens (JWT), and settle

payments automatically without issuing traditional, easily compromised API keys.[45] By integrating similar EIP-3009 gasless transaction standards and escrow mechanics directly into the FABIABox's modular services, the company transforms itself into the central clearinghouse of the autonomous workforce.[43] Every time an agent processes a document, authenticates a human user, or issues a legal signature via SquadSign, an automated micro-transaction is executed. This aligns perfectly with the revenue-sharing demand curve observed among developer-first builders, establishing a nonlinear, highly scalable revenue trajectory that is divorced from human headcount.[1]

Revenue Tier	Pricing Mechanism	Value Delivery	Target Buyer Profile	Economic Moat
Hardware (CapEx)	$5,000 unit sale MSRP	Sovereign silicon node (AMD/NVIDIA), immutable NixOS stack.	IT Procurement, CISOs, Compliance Managers.	Physical tangibility, localized execution, EU AI Act compliance.
MCP Services (OpEx)	M2M Micropayments (fractional cents per call)	LayerExtract (Data), DocuOstream (SquadSign legal execution).	Developers, Autonomous Agent Orchestrators.	SLA-backed reliability, EIP-3009 gasless settlement, JWT validation.

References

1. Enterprise AI trends in 2026: Sovereign, agentic, edge, AI factories - Spectro Cloud, accessed on June 29, 2026, https://www.spectrocloud.com/blog/enterprise-ai-2026-trends
2. AI Act State of Play – Key Obligations Postponed and Amended, Alongside New Guidance, accessed on June 29, 2026, https://www.skadden.com/insights/publications/2026/05/ai-act-state-of-play
3. EU Digital Identity Act - dewa, accessed on June 29, 2026, https://dewa-id.com/en/eidas
4. NVIDIA just announced NemoClaw at GTC, built on OpenClaw - Reddit, accessed on June 29, 2026, https://www.reddit.com/r/openclaw/comments/1rw05g5/nvidia_just_announced_nemoclaw_at_gtc_built_on/
5. Edge Computing Market Size & Share Report, 2026-2033 - Grand View Research,

accessed on June 29, 2026, https://www.grandviewresearch.com/industry-analysis/edge-computing-market
6. Edge Computing Market Size, Share, Industry Analysis - MarketsandMarkets, accessed on June 29, 2026, https://www.marketsandmarkets.com/Market-Reports/edge-computing-market-133384090.html
7. Edge AI Market Report 2026 - Research and Markets, accessed on June 29, 2026, https://www.researchandmarkets.com/reports/6226171/edge-ai-market-report
8. Inside NemoClaw An Architectural Analysis of NVIDIA's Enterprise Security Stack for Autonomous AI Agents - MK Science Set Publishers, accessed on June 29, 2026, https://mkscienceset.com/articles_file/523-_article1781758516.pdf
9. EU Approves Delays and Other Amendments to Certain EU AI Act Obligations: What Businesses Should Know - Morgan Lewis, accessed on June 29, 2026, https://www.morganlewis.com/pubs/2026/06/eu-approves-delays-and-other-amendments-to-certain-eu-ai-act-obligations-what-businesses-should-know
10. EU AI Act compliance hits in 47 days. Here's what it actually requires from AI agent builders, accessed on June 29, 2026, https://www.reddit.com/r/AI_Agents/comments/1u8betp/eu_ai_act_compliance_hits_in_47_days_heres_what/
11. What the EU AI Act Enforcement Delay Actually Means for Your Organization - A-LIGN, accessed on June 29, 2026, https://www.a-lign.com/articles/eu-ai-act-enforcement-delay
12. Automated regulatory compliance for AI systems in the security domain: The case of dual-use deployment. - Open Research Europe - European Union, accessed on June 29, 2026, https://open-research-europe.ec.europa.eu/articles/6-83
13. e-Wallet, EUDI Wallet, and what eIDAS 2.0 really changes - Cases & Insights, accessed on June 29, 2026, https://blog.e-boks.com/e-wallet-eudi-wallet-and-what-eidas-2.0-really-changes
14. eIDAS 2: Dates, approach, news and EUDI wallet - Tecalis, accessed on June 29, 2026, https://www.tecalis.com/blog/eidas2-eidas-2-regulation-eudi-20-europe-eu-wallet
15. Electronic Signatures vs. Digital Signatures: What's the Difference? - Proof, accessed on June 29, 2026, https://www.proof.com/blog/electronic-signatures-vs-digital-signatures-whats-the-difference
16. E-Signature & Digital Identity Laws – eIDAS 2.0 Becomes Reality - stp.one, accessed on June 29, 2026, https://www.stp.one/en/blog/de/blog/e-signature-digital-identity-laws-eidas-2.0-wird-realit%C3%A4t
17. Digital Identity in Notary Services: UK Essentials - Yousign, accessed on June 29, 2026, https://yousign.com/blog/digital-identity-notary-services
18. Electronic signatures in transactions - Noerr, accessed on June 29, 2026, https://www.noerr.com/en/insights/electronic-signatures-in-transactions

19. Personal AI Supercomputer Powered by Blackwell | NVIDIA DGX Spark, accessed on June 29, 2026, https://www.nvidia.com/en-us/products/workstations/dgx-spark/
20. NVIDIA DGX Spark from NVIDIA Corp. - Grace Blackwell GB10 brings local AI into the rack, accessed on June 29, 2026, https://www.ad-hoc-news.de/boerse/news/ueberblick/nvidia-dgx-spark-from-nvidia-corp-grace-blackwell-gb10-brings-local-ai/69649436
21. Nvidia unveils RTX Spark computer chip with up to 20 cores, RTX 5070 GPU and 128GB RAM - GSMArena.com news, accessed on June 29, 2026, https://www.gsmarena.com/nvidia_unveils_rtx_spark_computer_chip_with_up_to_20_cores_rtx_5070__128gb_ram-news-73061.php
22. AMD challenges Nvidia's DGX Spark with $3,999 Ryzen AI Halo with Windows 11 support — Strix Halo desktop undercuts Nvidia by $700, packs 128GB of unified memory, accessed on June 29, 2026, https://www.tomshardware.com/desktops/mini-pcs/amd-challenges-nvidias-dgx-spark-with-usd3-999-ryzen-ai-halo-with-windows-11-support-strix-halo-desktop-undercuts-nvidia-by-usd700-packs-128gb-of-unified-memory
23. Nvidia Unveils RTX Spark, an Arm-Based Superchip for Windows PCs, accessed on June 29, 2026, https://www.pcmag.com/news/nvidia-rtx-spark-reinvent-pc-computex-2026
24. AMD Ryzen™ AI MAX+ 395 Processor: Breakthrough AI Performance in Thin and Light, accessed on June 29, 2026, https://www.amd.com/en/blogs/2025/amd-ryzen-ai-max-395-processor-breakthrough-ai-.html
25. AMD Ryzen™ AI Halo for AI Developers, accessed on June 29, 2026, https://www.amd.com/en/products/processors/desktops/ryzen/ryzen-ai-halo.html
26. Framework Desktop Deep Dive: Ryzen AI Max - Blog, accessed on June 29, 2026, https://community.frame.work/t/framework-desktop-deep-dive-ryzen-ai-max/66002
27. AMD's CEO Destroyed NVIDIA's Most Expensive Supercomputers With a $1,500 Lunch-Box PC! - YouTube, accessed on June 29, 2026, https://www.youtube.com/watch?v=LgSI3WNOTP8
28. Deploy Gemma 3 on GPU Cloud: Complete Guide for All Variants | Spheron Blog, accessed on June 29, 2026, https://www.spheron.network/blog/deploy-gemma-3-gpu-cloud/
29. Gemma 3 VRAM Requirements — 1B, 4B, 12B, 27B GPU & Mac Guide (2026), accessed on June 29, 2026, https://willitrunai.com/blog/gemma-3-local-inference-guide
30. Gemma 3 - How to Run Guide | Unsloth Documentation, accessed on June 29, 2026, https://unsloth.ai/docs/models/tutorials/gemma-3-how-to-run-and-fine-tune
31. Gemma 3 Fine-tuning now in Unsloth - 1.6x faster with 60% less VRAM - Reddit, accessed on June 29, 2026, https://www.reddit.com/r/LocalLLaMA/comments/1jba8c1/gemma_3_finetuning_now_in_unsloth_16x_faster_with/

32. Using NixOS for Immutable Infrastructure and Declarative Configuration - GoCodeo, accessed on June 29, 2026, https://www.gocodeo.com/post/using-nixos-for-immutable-infrastructure-and-declarative-configuration
33. Code In, Cluster Out: Building Reproducible Edge Kubernetes with NixOS, K3s, and Forgejo, accessed on June 29, 2026, https://dev.to/ces0712/code-in-cluster-out-building-reproducible-edge-kubernetes-with-nixos-k3s-and-forgejo-i08
34. The Immutable Linux Paradox - Project Discussion - Ubuntu Discourse, accessed on June 29, 2026, https://discourse.ubuntu.com/t/the-immutable-linux-paradox/66456
35. Updating unreliable appliances - Deployments - NixOS Discourse, accessed on June 29, 2026, https://discourse.nixos.org/t/updating-unreliable-appliances/37782
36. Safer AI Agents & Assistants with OpenClaw | NVIDIA NemoClaw, accessed on June 29, 2026, https://www.nvidia.com/en-us/ai/nemoclaw/
37. NVIDIA NemoClaw, accessed on June 29, 2026, https://docs.nvidia.com/nemoclaw/user-guide/openclaw/home
38. A Deep Dive Into MCP and the Future of AI Tooling | Andreessen Horowitz, accessed on June 29, 2026, https://a16z.com/a-deep-dive-into-mcp-and-the-future-of-ai-tooling/
39. The Rise of MCP: Protocol Adoption in 2026 and Emerging Monetization Models - Medium, accessed on June 29, 2026, https://medium.com/mcp-server/the-rise-of-mcp-protocol-adoption-in-2026-and-emerging-monetization-models-cb03438e985c
40. Is building paid/premium MCP servers actually a viable business? Or am I missing something obvious? - Reddit, accessed on June 29, 2026, https://www.reddit.com/r/mcp/comments/1mbpjr1/is_building_paidpremium_mcp_servers_actually_a/
41. Documentation | TollBit, accessed on June 29, 2026, https://tollbit.com/docs/
42. How I built a payment gateway for my AI Agents to pay each other in USDC - Reddit, accessed on June 29, 2026, https://www.reddit.com/r/LangChain/comments/1t4tcfv/how_i_built_a_payment_gateway_for_my_ai_agents_to/
43. Sam Altman backs “micropayment” model for AI agents to compensate publishers, accessed on June 29, 2026, https://www.niemanlab.org/2026/05/sam-altman-backs-micropayment-model-for-ai-agents-to-compensate-publishers/
44. Your complete web stack for the agentic internet - TollBit, accessed on June 29, 2026, https://tollbit.com/mcp
45. gener8tor Luxembourg, accessed on June 29, 2026, https://www.gener8tor.com/investment-accelerators/luxembourg
46. Investment Accelerators at gener8tor, accessed on June 29, 2026, https://www.gener8tor.com/investment-accelerators
47. gener8tor | Accelerating the Best & Brightest, accessed on June 29, 2026, https://www.gener8tor.com/

48. Unit Sofia - European Laboratory for Learning and Intelligent Systems, accessed on June 29, 2026, https://ellis.eu/research/sites/unit-sofia
49. INSAIT | Institute for Computer Science, Artificial Intelligence and Technology., accessed on June 29, 2026, https://insait.ai/
50. Bulgaria launches AI research centre to woo science and tech talent, accessed on June 29, 2026, https://sciencebusiness.net/news/bulgaria-launches-ai-research-centre-woo-science-and-tech-talent
51. Bulgaria will have its own AI factory - a project of INSAIT and Sofia Tech Park for 90M EUR, accessed on June 29, 2026, https://insait.ai/bulgaria-will-have-its-own-ai-factory-a-project-for-90m-eur/
52. With Over 1000000 BGN in Donations, INSAIT Launches EXPLORER 2026 - a unique AI and Deep Tech bachelor program, accessed on June 29, 2026, https://insait.ai/with-over-1000000-bgn-in-donations-insait-launches-explorer-2026-a-unique-ai-and-deep-tech-bachelor-program/