Ahmed Elmalt

Principal Software Engineer

Principal Engineer with 20+ years architecting large-scale AI, distributed, and platform systems, currently lead PE for Alexa's communications organization. Architect of Alexa's real-time visual intelligence platform (patent filed) and of the multi-agent, media, and partner-integration platforms behind Alexa's LLM-era communications experiences. Sets multi-year technical strategy with VP and SVP leadership, converges fragmented architectures into shared platforms, and negotiates architecture directly with external partners at billion-user scale.

ahmed@eknowledger.com | +1 (804) 833-3171 | Irvine, CA | linkedin.com/in/eknowledger

KEY ACHIEVEMENTS

- Chief architect of Alexa's real-time visual intelligence platform for smart eyewear and multimodal devices, a net-new product category taken from concept to funded beta, with the architecture and multi-year strategy presented to Alexa VPs and the Amazon Devices SVP. **Patent filed.**
- Lead architect for bringing audio and video calling from one of the world's largest communication platforms onto the Alexa device fleet, authoring the strategy and counter-proposal that restored a paused program to a signed commercial agreement.
- Rebuilt Amazon's real-time media platform so any engineer, not only media specialists, can ship production real-time media and voice-agent applications: a working application went from over **1,100 lines to 8**, and from weeks to **under an hour**.
- Founding member of Alexa's monetization engineering organization, architecting the advertising platform foundation that reached **$132MM ARR (78% YoY)** with a **$700MM+** projected run rate.
- Architected Alexa's first contextual targeting platform, a **$1B+** projected-revenue system at 100+ TPS and 99.9% availability, delivering **3x engagement lift** over existing placements.
- Published researcher in conversational reinforcement learning, and Alexa API Bar Raiser.

EXPERIENCE

Principal Software Engineer

May 2024 – Present

Amazon

Irvine, CA

Alexa Communications

Lead Principal Engineer for Alexa's communications organization: 3 pillars, 17 teams, 150+ engineers, and 107+ cloud services, owning platform architecture and multi-year technical strategy across real-time multimodal AI, agentic communication systems, third-party VoIP partner integration, and real-time media infrastructure.

Real-Time Visual Intelligence Platform

Alexa's platform for real-time visual understanding across smart eyewear, multimodal screen devices, and mobile.

- Owned architecture for a net-new product category from concept through funding to beta. **Presented the architecture and multi-year technical strategy to Alexa VPs and the Amazon Devices SVP**, aligning senior leadership on platform direction and advising the joint hardware, software, and modeling roadmap.
- Established the platform's two core abstractions: a normalized enrichment as the universal output contract, and pluggable visual processors spanning device to cloud. Separated orchestration control plane from frame data plane so each scales independently, and authored the cloud-device capture protocol now used as the partner integration contract.
- Invented a speculative capture-gating architecture split across the device-cloud boundary (patent filed). An on-device classifier tuned for recall captures the transient moment; a cloud classifier tuned for precision, holding interaction history, drops unintended frames pre-persistence.
- Designed hierarchical visual memory implementing retrieval-augmented generation over video, fusing captions, embeddings, and interaction history at index time, which lifted multi-hour visual recall **from 15% to 36%** while keeping the layer model-agnostic.
- Bootstrapped the visual memory service myself to de-risk the design before staffing: wrote the service, model, and CDK infrastructure packages, stood up the build pipeline and deployment, and prototyped the key-frame selection algorithm that cut context-window consumption on the ingest path.
- Ran the retrieval-quality and model experiments that set the science direction, diagnosing a caption-to-query semantic gap as the ceiling on recall and prototyping structured dense captioning against it. Benchmarked frontier streaming models on egocentric shopping footage and against public online-video benchmarks to establish what was and was not achievable at launch.
- Extended the interaction model past spoken and written input to proactive, non-verbal requests from gesture, action, and scene change, and opened the platform to partner domain teams through one versioned processor contract, secured by a four-tier trust architecture that enforces data isolation structurally rather than by policy.
- Built the cost model that redirected the optimization program and device roadmap: the dominant cost is per-frame model invocation, not reasoning inference, making key-frame selection and on-device gating the primary levers. Also API bar raiser on public camera specs.

Third-Party VoIP Integration with a Global Communication Platform

Architecture and hard-problem engineering to bring audio and video calling from one of the world's largest communication platforms (billions of monthly users) onto the Alexa device fleet, across 40+ device variants and three operating systems under a signed agreement.

- Established write-once-compile-everywhere as the architectural thesis, replacing N-by-M per-OS glue with one ABI-stable interface-definition contract compiled to per-OS projections, over a layered runtime separating partner-owned media and signaling from Amazon-owned OS primitives, a reusable communications SDK, and a per-OS container app. Now the committed architecture, proven by end-to-end calling.
- Ran the architectural down-select that set program direction, evaluating embedded browser-engine, Bluetooth-profile, cloud-streaming, and WebAssembly alternatives against device-fleet reach, memory envelope, and real-time viability. Each reached only a small fraction of the installed base, leaving native execution as the only viable path.

- Proved the architecture hands-on rather than on paper: cross-compiled the third-party toolchain onto the constrained device target to demonstrate the build path was real, read the contact-resolution service source to verify the pipeline reuse the plan depended on, and audited the calling proof-of-concept code to separate what was genuinely working from what was stubbed.
- Redefined the resource model for third-party code on constrained consumer hardware, correcting the governing memory budget and an inverted failure mode in which a media-intensive call degraded the assistant rather than the call. The resulting isolation and admission model was adopted into the platform roadmap.
- Re-architected contact resolution to reconcile the partner's strict on-device privacy posture with Alexa's cloud-first architecture, replacing bulk address-book upload with progressive call-time resolution so only contacts actually called ever leave the device, at equivalent steady-state latency. Partner aligned on the design.
- Designed identity and account linking for personal accounts on communal, frequently screenless devices using tiered ownership verification and delegated approval, and defined the cross-OS real-time media contract that unblocked three operating systems: one buffer-ownership and data-flow policy spanning audio and video, plus echo-cancellation and hardware-codec capability negotiation across heterogeneous device audio hardware.
- Authored the Director and VP platform assessment diagnosing why deep third-party integration was structurally unsustainable, and recommended one consolidated third-party application platform over disconnected fixes. **Advised Alexa and Devices VPs on the resulting technical strategy**, then authored the counter-proposal that answered the partner's re-engagement conditions: **partner aligned, term sheet signed, program resumed as a top-level company goal**.
- Led three multi-day joint engineering offsites with the partner's platform and international organizations, and founded the standing weekly working group operating a normative decision log, now the organization's pattern for cross-company governance.

Real-Time Media Platform and Agentic Toolchain

A Python-first SDK, engine, runtime, and toolchain for Amazon's real-time media cloud: the signaling and media backbone for every real-time communication on Alexa, where session orchestration serves 43% of Alexa+ traffic and the media relay carries 100% of Alexa calling traffic (including ~10% from Ring and Smart Home) across 42 production services.

- Took over a specialist-only platform and rewrote it end to end as principal author, replacing a four-language runtime chain with a direct Python-to-pipeline path on the same production runtime and codecs, collapsing a working media app to eight lines and under an hour.
- Personally shipped **200+ code reviews and 650+ commits** across the nine packages I own, roughly **200K lines** backed by **~1,200 automated tests**, spanning the engine, SDK, CLI, plugin system, transport and signaling layers, container runtime, and release pipeline.
- Set the language boundary as the foundational architectural decision: a language-agnostic engine on the native media C ABI with SDKs as replaceable wrappers, keeping the engine reusable and further language SDKs additive rather than a rewrite.
- Made a managed language defensible on a real-time media path through mechanism, not optimism: the media hot path stays in native threads with the interpreter lock released, so frames never queue behind application code, with garbage-collection pauses eliminated and cost surfaced only where an author opts in. Tuned to a **50ms real-time target** with a P99 regression gate.
- Replaced a transport-specific abstraction with a general typed graph model, holding the primitive set closed by contract so new capability arrives as patterns above it. That reframe carried the platform beyond calling into voice agents, live view, transcription, and ingest.
- Wrote the pattern layer that eliminated per-application state machines: declarative topology separating the once-built shared core from per-participant legs, participant role negotiated in signaling, and swappable layout. Solved multi-party dynamic reconfiguration so participants join and leave branch-locally against a core that never leaves the playing state.
- Built the developer toolchain as platform architecture, not convenience: a production-equivalent container reachable in one command on any OS, template-driven scaffolding, four progressive extension levels, and a command surface usable non-interactively so agents drive the same interface as engineers, backed by a no-scenario-no-ship verification contract.
- Built the team's agentic engineering layer on the principle that the machine interface is the primary API: an agent package covering specification authoring through review, publication, and code review, generating surfaces for three agent runtimes from one source, with an audit framework that caught real defects pre-review. Multiple partner teams now build on it in production, on infrastructure underpinning 180M+ weekly real-time sessions across communications, Smart Home, and Visual ID.

Agentic Incoming Communications Platform

A multi-agent AI platform that ingests, reasons over, and acts on a customer's inbound communications across channels, scaling toward 100K+ users and 1M+ messages per day.

- Founded and led the engineering and science working group that became the program's decision forum, and authored the technical strategy and roadmap taking it from a single-channel email prototype to a multi-channel orchestration platform.
- Converted the full requirement surface into five architectural themes with explicit one-way-door decisions and published design tenets, making binding constraints legible to leadership and stopping the MVP from foreclosing the multi-channel vision.
- Architected the multi-agent orchestration layer: a supervisor agent performing task planning and decomposition, domain agents communicating only through an event bus to preserve loose coupling, a shared task store for context and history, and hierarchical tasks with dependency waits, event-driven handoffs, and rescheduling, replacing brittle manual pipelines with a horizontally scaling agent ecosystem.
- Built the early prototypes and proofs-of-concept that validated the architecture under scale before the team committed to it, working hands-on through the agent orchestration path, the data pipeline, and the generation models rather than handing off a design document.
- Designed the channel abstraction letting one intelligence layer serve email today and messaging, SMS, and voice later without rewrite, through a universal message schema, channel adapters normalizing protocol and capability differences, and a plugin architecture for independent channel development.
- Held the platform to **sub-2-second latency** and **>85% classification accuracy** while coordinating multiple model invocations per message, using parallel and sequential orchestration stages, context propagation, and circuit-breaking for graceful degradation on inference failure.
- Architected and helped implement the synthetic data generation pipeline that unblocked model development under privacy constraints, using persona-driven generation, inbox grounding, and topic distribution modeling to build representative corpora without real customer data.

- Closed a validation gap that would have carried a misleading accuracy signal into production, replacing a circular synthetic-to-synthetic evaluation loop with a corrective framework covering sample sizing, experimental controls and cross-validation, model selection, phased real-world validation, and a customer-feedback learning loop.

Alexa+ Communications Platform Architecture

- Led the technical strategy and architecture for **Communicate with Contacts**, a multi-organization S-Team initiative enabling zero-friction communication across SMS, email, Bluetooth, and RCS on Alexa's next-generation LLM runtime, integrating natural language understanding, dialog management, device software, and cloud services across seven Alexa and AGI organizations.
- Developed a prompt architecture combining chain-of-thought reasoning, finite-state machine structure, and runtime model selection between large and small models, cutting input tokens by **over 40%** and user-perceived latency by **60%** while holding accuracy above **98%**.
- Primary technical lead for communications experiences through the **Alexa+ launch**, shipping seven net-new communication experiences and five new experts across a portfolio serving 33M monthly and 4.2M daily active accounts. Co-led the messaging expert refactoring that reduced inference calls and latency, and designed the structured state-machine prompt architecture that language engineering extended to lift accuracy and CSAT.
- Designed a three-track speculative execution architecture that attacked multi-second messaging latency across millions of daily requests. Speculatively resolving contacts and pre-generating content ahead of the request removed several hundred milliseconds to over three seconds of delay, delivered as reusable infrastructure across channels.
- Set the multi-year strategy converging the organization's fragmenting client architecture, replacing per-platform calling implementations with one write-once architecture on portable UI primitives and capability-driven per-device rendering, and decoupled the design system from the rendering runtime so each migrates independently.
- Set technical direction on organization-wide questions including agentic orchestration convergence, contact resolution modernization through applied ML, service consolidation and fragmentation reduction, carrier calling sunset, and customer data handling and deletion compliance across international markets and child-data contexts.
- Drove AI-assisted engineering adoption across the organization **from 30% to 85% in a year**, authoring the guidance on evolving code review for AI-accelerated development, building the agentic specification and review tooling the team now runs on, and writing the majority of my own designs as working code through agentic development.

Principal Software Engineer

Mar 2023 – May 2024

Amazon

Irvine, CA

Alexa Monetization

Founding member of Alexa's monetization engineering organization, architecting the advertising platform foundation behind Alexa's first contextually and behaviorally targeted advertising.

- Led cross-organizational technical strategy across Alexa Info, AI, Devices, Proactive, and Amazon Ads, resulting in **SVP approval of a three-year technical strategy**, dedicated funding across organizations, and the launch of Alexa's first behaviorally targeted advertising.
- Architected and built Alexa's conversational advertising platform, delivering interactive audio and display advertising with LLM-driven interaction, and implemented the context-grounding system that made model behavior deterministic through semantic prompt construction and exemplar-based fine-tuning.
- Implemented chain-of-thought prompt chaining to execute deterministic advertising workflows, and built the real-time attribution service correlating millions of daily events across voice and touch interfaces at **99.99% accuracy** using event correlation and temporal heuristics.
- Architected the **Alexa Audiences Platform**, transforming petabyte-scale interaction data into interest, demographic, and in-market audiences through a federated serverless architecture pairing a data lake with distributed compute, and built the extensible orchestration layer driving audience computation from signal ingestion to privacy-compliant publishing.
- Built the ML pipelines converting raw interactions into standardized interest and demographic signals, plus the integration layer publishing Alexa audiences into Amazon's advertising platforms, and implemented privacy-first processing throughout: pseudonymization, retention, consent, underage exclusion, opt-outs, and data lineage.
- Built extensible APIs and SDK components integrating advertising into existing application frameworks against industry viewability and tracking standards, leading 10+ engineering teams through state management, navigation, and overlay tracking.

Senior Software Engineer

Jun 2018 – Mar 2023

Amazon

Irvine, CA

Alexa Monetization: Contextual Targeting & Audio Advertising

Founding engineering work on Alexa's advertising targeting platform.

- Architected and led development of Alexa's first contextual targeting platform, a mission-critical system with **$1B+ projected revenue**, scaled to 100+ TPS at P50/P90 latencies of 75ms/150ms and **99.9% availability** across multiple ad products and content providers.
- Engineered the ML inference architecture on GPU and ECS sidecar containers with host-level IPC to hold **sub-100ms latency**, with hot model reload, multi-AZ deployment, and drift detection, over an end-to-end pipeline spanning distributed SageMaker training, 100+ contextual features, real-time inference, and monitoring.
- Engineered the multi-stage signal generation and ranking system combining BERT semantic scoring, semantic graph reasoning over OWL/RDF/SHACL ontologies, and intent classification, evaluated in real time against retail catalog signals with dynamic confidence scoring, delivering **3x higher engagement (1.8%+ CTR)**.
- Implemented the production graph-traversal algorithms pairing breadth-first search with semantic reasoning to generate and score targeting signals over Amazon's Alexa knowledge graph, and extended it into a unified schema bridging Alexa with the e-commerce ontology.
- Implemented a configuration-driven DAG orchestration engine with dynamic branching and parallel execution, cutting new targeting strategy delivery from weeks to days, and created the self-service publisher platform and visualization tools behind cross-domain onboarding.

- Architected and implemented the audio advertising insertion and tracking system for Alexa news on AWS Lambda, ElastiCache, and SNS, serving millions of daily requests at **P99 150ms**, and wrote the new application protocol directives, event handlers, and content-focus extensions that keep playback continuous through ad insertion, skip, and focus transfer.

Alexa Skills Store: Search & Conversational Discovery

Discovery systems for Alexa's third-party skill ecosystem.

- Led design and implementation of a unified search and discovery system integrating third-party skills, first-party features, and voice command catalogs into one conversational experience, with a search arbitration layer that filters, ranks, and merges across catalogs while preserving relevance ordering and heterogeneous launch protocols at **P90 sub-250ms**.
- Implemented the scalable data provider architecture in Java on AWS Lambda, handling complex metadata relationships and real-time content hydration across millions of catalog entries, plus a fault-tolerant integration layer with filtering, tag-based categorization, and ranking.
- Built the conversation engine driving natural discovery flows, with presentation strategies for mixed content types and contextual recommendation across Alexa surfaces.
- Architected a conversational skill discovery agent using deep reinforcement learning, training a Deep Q-Network over 150,000 steps and 3,000 dialog episodes to run multi-turn dialogs. Raised discovery success **from 73.4% to 77.0%** while shortening average dialog length ($p<0.0001$) and published the work (arXiv:1911.06747).
- Built and deployed the production A/B testing framework that validated the policy gains, and delivered the serverless architecture spanning **150K+ skills across 190+ categories**, achieving a **44% conversion improvement** with **75% fewer critical incidents** versus the prior system.

Director of Engineering

Oct 2015 – Jun 2018

GoFormz Inc.

San Diego, CA

Employee #9 at a digital transformation startup; built the engineering organization and platform through 100+ enterprise clients, 1M+ users, and two funding rounds totaling $10M+.

- Built the engineering team from the ground up to 25+ engineers, owning budget and infrastructure cost, reporting strategy monthly to the CEO and board, and establishing the CI/CD pipeline and DevOps practice that **cut deployment time by 60%** while improving release stability.
- Led the architecture and migration of a monolith into a scalable service-oriented platform using domain-driven design, enabling **100K+ paper form digitizations annually**, and built the distributed capture and transformation pipeline on event-driven architecture and CQRS.
- Led design and implementation of the company's internal microservice platform in C#/.NET, including a resilient execution framework, API rate limiting, and event-driven communication over event sourcing, so product teams could stand up services rapidly.
- Built the third-party application runtime, SDK, and marketplace behind **hundreds of integrations** with major SaaS providers, implementing a cluster-management algorithm balancing load across nodes by availability and capacity through node election and service-bus messaging.
- Implemented the runtime resiliency layer for long-running third-party workloads, handling failure, retry, timeout, circuit-breaker fail-fast for corrupt commands, and bulkhead capacity isolation, with non-blocking metric collection on .NET TPL and Reactive Extensions.
- Engineered the platform security foundation: an identity microservice supporting OAuth2/OpenID and SAML2 SSO with enterprise provider integration, RBAC authorization cached in Redis, and an internal API gateway handling authentication, authorization, auditing, and distributed transactions via sagas.
- Designed and helped implement the workflow engine letting third-party developers define workflows over captured data, covering transformation, application integration, control flow, and HTTP operations, over C# REST services with MongoDB and Azure Storage.

Principal Engineer

Oct 2012 – Oct 2015

Cornerstone OnDemand

Santa Monica, CA

Key member of the small architecture group behind Extend, a PaaS enabling clients, partners, and third-party developers to build and monetize applications on a talent management suite serving 10K+ clients.

- Designed the platform's application development language and wrote the DSL compiler emitting intermediate platform instructions that coordinate application actions across cloud services.
- Co-designed the platform execution intermediate language and built the runtime virtual machine that executes and coordinates those instructions across the service estate.
- Designed a second internal DSL for data transformation, expressing rules in near-natural language, and implemented its interpreter, compiling rules into optimized units of work for the runtime.
- Built the rules engine service in C# over WCF and MSMQ, generating execution graphs from DSL rules to maximize resource utilization during bulk data loading.
- Implemented the platform data services engine translating OData 4.0 queries into in-memory object graphs via C# expression trees and compiled queries for an optimized retrieval pipeline.
- Contributed to the developer portal, delivering a cloud-hosted IDE on AngularJS with C# backend services so third-party developers could author, debug, and ship application code against the compiler.
- Owned core platform services including language services, runtime services, and cluster management, and built the data transformation wizard clients used to author and schedule rules.

Lead Engineer

Oct 2010 – Oct 2012

General Electric (GE Healthcare)

Richmond, VA

Next-generation hospital smart room and staff compliance platforms, deployed across hundreds of medical facilities worldwide.

- Architected and implemented a scalable real-time location server supporting multiple tracking technologies (RF, infrared, ultrasound) and multiple hardware vendors, serving as the central hub for tracking people and assets across medical facilities.

- Implemented a pluggable sensor adaptor architecture over a pub/sub model on WCF netTCP, socket APIs, and ActiveMQ, scaling to millions of beacon events per hour with parsing, transformation, and distribution to consuming applications.
- Designed and implemented the indoor positioning algorithms fusing RF, infrared, and ultrasound through triangulation and proximity techniques, achieving **under 2% error** and cutting equipment search time **from 24 minutes to 4**.
- Led architecture and development of a joint solution with a major networking vendor unifying wireless network and context-aware hospital software, bridging WiFi sensory data and real-time location systems for the first time into one facility tracking solution.
- Led architecture of a hand hygiene compliance platform combining RF and infrared signals with probabilistic proximity detection, and co-implemented the core algorithm computing missed compliance opportunities in real time.
- Partnered with the chief hardware engineer to build the device-to-server wireless transmission protocols in C# over socket APIs, WCF, and ActiveMQ, and built the compliance dashboards on C#, Reactive Extensions, SQL Server, and SSRS used across clinical and administrative groups.

Earlier Experience

2003 – 2009
Cairo, Egypt

Lead Engineer / Senior Software Engineer, Bluebridge Technologies (2006 – 2009)
Led architecture and delivery of enterprise solutions for North American and European clients, including a business workflow and resource management platform for Austria's largest telecommunications company on C#, SharePoint, WCF, MSMQ, and BizTalk. Contributed to **OmniPage v16**, the world's leading OCR software, where I designed and wrote the interpreter that translates the OCR engine's binary output into Open XML across .NET, C#, C++/CLI, and native C++, and led an OCR workflow automation framework built with Microsoft. Authored the agile secure SDLC framework adopted company-wide.

Senior Software Engineer, Arcom (2005 – 2006)
Led architecture and development of Saudi Arabia's e-government portal on C#, ASP.NET, and SQL Server, serving **30M+ citizens and 100M+ monthly transactions**, and engineered a high-security enterprise information monitoring system for a national ministry on .NET Remoting, COM+, and WSE, processing millions of daily sensitive transactions, plus e-learning platforms for national defense institutions under CMMI.

Software Engineer, Bayan-Tech (2003 – 2005)
Built web portal and e-learning modules on C#, ASP.NET, and SQL Server, engineered XEdit, an XML-based e-learning content editor for mobile devices, and wrote reusable infrastructure components including data providers, error logging, and globalization controls.

PUBLICATIONS & PATENTS

- **Patent pending, first-named inventor of five.** Visual context integration with natural language reasoning: speculative visual capture gating across a device-cloud boundary using paired high-recall and high-precision classifiers. Filed 2026; details confidential pending publication.
- **Towards Personalized Dialog Policies for Conversational Skill Discovery** (arXiv:1911.06747). Deep Q-Network agent outperforming a rule-based policy on success rate and dialog length with significance.
- Authored the advertising-targeting architecture whitepapers used for Director and VP strategic planning across Alexa and Amazon Ads, plus multi-year architectural plans shaping organization-wide approaches to data handling and targeting.

TECHNICAL LEADERSHIP

- **Alexa API Bar Raiser.** Review and approve public API specifications across camera, contacts, and communications domains.
- **Working groups founded and led.** Four cross-organizational technical forums across visual intelligence, partner integration, agentic communications, and client architecture convergence, including a cross-company forum operating a normative decision log.
- **Mentorship and hiring.** Mentored 12+ engineers through formal programs, developed senior engineers into architecture owners contributing to multiple Principal and senior promotions, and served on Principal Engineer, senior product, and TPM loops.
- **Podcast founder and host.** Kalam Falsfa, the leading philosophy podcast in the Middle East, and co-founder of Ask Developer, a technology podcast for the Middle Eastern developer community. Speaker on microservice architecture, compiler and DSL design, and high-performance runtimes at company symposia.

SKILLS

AI & Machine Learning: Large language models, agentic and multi-agent architectures, MCP agent tooling, retrieval-augmented generation, vector search and embeddings, reinforcement learning (DQN), transformer models, prompt architecture, ML pipelines, and model evaluation.

Platform & Distributed Systems: Platform and API design, multi-tenant extensibility, SDK and developer experience, plugin architectures, DSL and compiler design, microservices, event-driven architecture, CQRS and event sourcing, domain-driven design. Languages: Python, Java, C#/.NET, TypeScript, and C/C++.

Real-Time & Device: WebRTC, GStreamer, VoIP and telephony, real-time media pipelines, low-latency inference, cross-platform native integration, embedded and resource-constrained systems, IoT and RTLS.

Cloud, Data & Security: AWS and Azure, serverless architecture, containerization, CI/CD, petabyte-scale data processing, DynamoDB, PostgreSQL, Redis, Elasticsearch, identity and access management (OAuth2, OpenID, SAML2), privacy-by-design, GDPR/CCPA, capability-based authorization.

EDUCATION

M.S. Computer Science, Maharishi University of Management, Fairfield, IA · *2009 – 2012*
B.S. Computer Science, Menoufia University, Egypt · *2004*
Amazon Machine Learning University, Graduate · *2019*
Project Management (PMP) Diploma, RITI, Egypt · *2008*

Last update June, 2026