# AI Architect Certification Program Brochure

## Program Overview

This two-day training program prepares technology managers, software engineers, solution designers, product owners, risk managers, compliance officers, and artificial intelligence support experts to design, secure, evaluate, and scale production-grade enterprise artificial intelligence platforms. Grounded in artificial intelligence system design principles, the curriculum focuses on concrete engineering infrastructure, deterministic control boundaries, unit economics, and governance stacks required for mission-critical deployment.

The program directly addresses the architectural decisions needed to plan developments and deployments while balancing technical complexity, operational cost, and organizational risk. The program uses a progressive, modular pedagogical approach for maximum flexibility. It equips cross-functional managers and experts with actionable architecture recommendations, decision frameworks, and delivery plans.

A core advantage of the program is its fully vendor-agnostic approach. It covers common execution patterns, primitive layers, and control capabilities without limiting focus to single ecosystems such as Anthropic, OpenAI, or Google. Because the curriculum concentrates on solution planning, architectural topology, and engineering options, there are no mandatory coding requirements linked to any specific programming language, platform or software framework. Our discussions are not tied to one vendor's tooling.

AI pilots stall in one meeting, when someone asks what one transaction costs, who approves the risky ones, and where the audit trail is. In two days you build the answers with architects, product owners, security managers and auditors in the room. You leave with effective AI recommendations you can use on Monday.

## Job Market Needs

The market demand for AI architects stems from a sharp shift in enterprise priorities: companies are moving past isolated prototypes and need architectures that embed AI directly into core business processes. Building a reliable system requires far more than picking a flagship model. It demands connecting probabilistic outputs to enterprise resource planning platforms, legacy databases, and operational workflows. Organizations need professionals who can take unmanaged experiments and turn them into resilient, monitored, and cost-controlled production platforms.

To bridge the gap between business objectives and technical execution, architects focus heavily on system topology, work decomposition, and trade-off management. They systematically divide workflows into AI-assisted tasks, deterministic software logic, and human oversight checkpoints. Their core responsibility is designing the end-to-end integration layer, setting up tool-calling mechanics, Model Context Protocol servers, and agentic state machines. Balancing model accuracy, execution latency, and unit economics ensures that automated execution delivers measurable business value without runaway token costs.

On the operational side, companies require risk and vulnerability management, security boundaries, and financial controls built into the platform from day one. AI architects implement fail-closed controls, access controls, prompt injection defenses, and immutable audit logging to meet compliance and security standards. By applying prompt caching, dynamic reasoning toggles, and context loading strategies, they keep inference costs predictable and ensure multi-step agentic workflows run safely at scale.

Meeting these market demands requires a specific set of technical, architectural, and governance skills that map directly to modern enterprise standards. Practitioners must master cloud delivery routes, integration protocol stacks, programmatic API mechanics, and vector context retrieval to manage system memory effectively. They need hands-on expertise in MLOps and LLMOps, including session correlation tracking, schema validation, and multi-tier evaluation ladders using model-as-a-judge audits. Finally, they must command governance and cost engineering frameworks, using architecture decision records, unit economics modeling, and board-level defense strategies to align technical deployments with business risk limits.

## Target Audience and Value Delivered

The program delivers targeted value across technical, operational, and executive roles by addressing their core system design and governance challenges:

- **Software Engineers and Architects:** Controlling the core integration layer, moving from basic interface prototyping to production-grade deployment using enterprise primitives, model context protocol servers, and agentic state machines.
- **Solution Designers:** Apply structured work decomposition to map complex workflows, separating probabilistic model tasks from deterministic software rules and human operational checkpoints.
- **AI Product Owners:** Control operational expenditure using unit economics, token throughput management, prompt caching strategies, and total cost of ownership modeling.
- **Chief AI Officers:** Obtain technical governance frameworks, control structures, and board-level defense strategies to transition organizational experiments into audit-ready enterprise platforms.
- **Software Testers and QA Engineers:** Build rigorous golden evaluation sets and implement multi-tier quality grading pipelines combining schema validation, model-as-a-judge audits, and production observability.
- **Chief Technology Officers:** Establish institutional memory through architecture decision records and isolate system risk using three-tier architectural boundaries across major cloud platforms.
- **MLOps and Platform Reliability Engineers:** Track non-deterministic failure modes in production, monitoring schema breaches, tool execution failures, token latency, and operational cost thresholds.
- **Compliance and Risk Managers:** Implement fail-closed guardrails, design stakes-based human review routing, and establish control layer ownership across model providers and internal platforms.

- **Enterprise Data Engineers and Database Architects:** Optimize context engineering, semantic retrieval pipelines, and working memory architecture to balance retrieval accuracy against context window costs.
- **Chief Information Security Officers and Security Engineers:** Deploy defensive controls against artificial intelligence threats, including direct and indirect prompt injection, privilege escalation, and data exfiltration.
- **IT Auditors and Enterprise Compliance Officers:** Inspect deterministic execution logs, inspect system deliverables, and evaluate immutable audit trails to verify AI technical and regulatory compliance.

The curriculum also benefits **all professionals** building custom agents and internal automations for individual or departmental use. It provides practical methods to structure prompts, handle data securely, and prevent common execution errors in everyday workflows. These principles help teams create reliable, cost-effective tools that streamline daily tasks without introducing operational risk.

## Curriculum Highlights

The Enterprise AI Architect certification program offers a curriculum focusing on practical system design, integration primitives, and governance engineering to plan, build, and scale production-grade AI platforms. You will gain technical practices and implementation approaches to:

- Develop an understanding of architectural isolation and work decomposition, learning to translate business objectives and requirements into secure, cost-effective, and resilient technical platforms.
- Implement defense-in-depth security using deterministic code hooks and fail-closed guardrails that physically block known paths ata exfiltration and prompt injection threats.
- Assess and design core building blocks including tools, model context protocol servers, subagents, and workflow orchestration to connect models safely to enterprise databases, cloud platforms, and operational systems.
- Model unit economics, manage total cost of ownership, route tasks across model tiering strategies, and apply prompt caching mechanics to lower execution latency.
- Dive into retrieval frameworks, vector indexing, and working memory architectures essential for grounding models on trusted internal records and preventing system hallucinations.
- Examine multi-tier evaluation ladders, schema validation checkers, session telemetry tracing, and versioned rollback strategies to monitor system health and correct operational drift over time.

This curriculum is structured to deliver hands-on learning through case-led testing, group discussions, and situational analysis, providing a strong blend of theoretical knowledge and practical insights.

## Learning Outcomes

Upon completion, you will be equipped to business needs with technical execution, transforming unpredictable probabilistic models into cost-effective, secure, and resilient enterprise software platforms. Key outcomes include:

- Use AI system design best practices and architectural isolation, building fault-tolerant validation software, memory management layers, and error recovery workflows rather than relying solely on foundation model selection.
- Enforce deterministic governance and multi-layered security controls using code hooks and automated quality gates that physically block privacy breaches and security threats rather than depending on prompt instructions alone.
- Orchestrate enterprise integration primitives and protocol stacks, utilizing tools, Model Context Protocol servers, subagents, and skills to connect models safely to enterprise databases, cloud pipelines, and operational applications.
- Engineer token unit economics and total cost of ownership by matching task complexity to model tiering strategies, configuring prompt caching, and setting enforceable spending limits across user workflows.
- Ground system outputs through authoritative retrieval frameworks and direct application programming interface tool calls, replacing model speculation and hallucinations without requiring costly model retraining.
- Implement automated quality gates and production observability, validating model outputs against software schemas before downstream execution while setting up telemetry tracing and rollback paths to correct operational drift.

This program equips you with tools to evaluate architectural options, establish control boundaries, and deploy enterprise artificial intelligence platforms within a robust governance structure.

This program offers a practical, framework-agnostic approach that prioritizes real-world system design over high-level theory. We deliver direct, actionable insights backed by experienced facilitators who build and audit enterprise systems. Tailored to meet the needs of architects, engineers, product owners, and governance specialists, the course covers the full spectrum of technical solution design, unit economics, and risk control. Aligning with emerging industry benchmarks, technical security guidelines, and compliance standards, this program ensures you are prepared to build reliable, audit-ready platforms while leading your organization into a scalable automated future..

## Hands-on Learning Experience

Our program emphasizes applied scenario analysis and rigorous business case discussions. Every participant receives access to decision frameworks, comparative design recommendations, and architectural reflection methodology. Real-world discovery dialogues and post-launch case audits enable you to evaluate complex trade-offs immediately, establishing a solid foundation to lead AI initiatives from a business and technical perspective. Pre-reading materials prepare you with shared terminology, reflection questions, and comparative design scenarios before classroom sessions begin. Facilitators with enterprise experience guide in-depth discussions on system design options, cost trade-offs, and operational impacts.

## Real-World, Impact-Driven Learning

This program grounds every technical concept in real-world enterprise realities, moving past idealized vendor demos to focus on how AI systems actually behave, and fail, in production. The curriculum teaches a disciplined system design approach that surrounds probabilistic models with deterministic software controls, effective work decomposition, and multi-layered security hooks. Through side-by-side comparative design setups, actual failure audits, and practical exercises in token economics, participants learn to build resilient architectures that isolate risk, enforce compliance, and optimize unit costs. By applying core integration primitives, context engineering strategies, and automated quality gates, technical leaders and architects gain the exact framework needed to turn unpredictable language models into audit-ready, cost-effective enterprise platforms.

## Enhancing Your Executive Profile and Networks

Completing the AI Architect certification marks you as a skilled practitioner expert in AI system design, technology governance, and platform architecture. Moving beyond certifications limited to particular model providers, this program equips you with the exact technical fluency and business frameworks needed for high-demand roles like enterprise AI architect, principal solution designer, AI product owner, and AI risk manager. Connecting engineering decisions to business accountability helps you step into high-stakes technology roles. You gain the confidence to manage unit economics and defend your architectural choices directly in front of product owners, executive committees and auditors.

Beyond individual skill building, the program creates an active, ongoing network of technology managers, system architects, and governance specialists. Through interactive discovery dialogues, cross-industry case analyses, and collaborative exercises on architecture decisions, you engage in candid peer debates with leaders from finance, healthcare, legal, and logistics. By working through shared failure modes, such as unmanaged token cost surges, compliance drift, and retrieval errors, you exchange field-tested strategies and benchmark your organization's technical maturity against industry peers. You leave with lasting professional connections and access to an executive alumni network focused on scaling trusted AI platforms.

## Admission Requirements

This certification is tailored for engineers, testers, product owners and managers, with an admission requirement of at least 5 years of experience in a relevant field, including technology, software engineering,

product management, or business functions. No coding background is necessary, as the program focuses on AI's managerial and strategic implications.

## About the Institutions

This certification is offered by **Copenhagen Compliance**, a globally respected leader in governance, risk, and compliance (GRC) implementation and training. As a prominent GRC management agency since 2014, Copenhagen Compliance understands the complexity of regulatory challenges that today's multinational organizations face. We believe that simply layering on more compliance is not the answer; instead, they equip professionals with deep-rooted solutions that effectively address the real issues organizations confront in governance, risk management, and compliance.

Copenhagen Compliance provides a comprehensive suite of GRC tools, software, and services that streamline these complex processes. Our global GRC platform includes standardized frameworks and guidance that help organizations achieve more efficient, automated GRC operations. From risk assessment and policy management to incident tracking, internal controls, and audit functions, Copenhagen Compliance offers a full spectrum of resources tailored to meet evolving regulatory demands and ensure a smooth, structured compliance workflow.

In collaboration with the **Information Security Institute**, Copenhagen Compliance further enhances the certification experience by bringing together a rich network of events, educational resources, and IA, IT and cybersecurity industry insights. The Information Security Institute supports safe technology adoption in global institutions, offering specialized modules that address critical regulations and frameworks, including AI, GDPR, cybersecurity and technology standards, ISO 42001, ISO 27001, SOX, and NIST, as well as corporate social responsibility and environmental, social, and governance requirements.

Together, Copenhagen Compliance and the Information Security Institute ensure that certification holders gain both foundational knowledge and cutting-edge insights into AI governance, risk, and compliance, well-prepared to address the regulatory, safety and business needs.

## Industry Recognition

Our certification is supported by partnerships with leading AI organizations and prominent academic institutions. With over 5,000 professionals certified by Copenhagen Compliance in privacy, data governance, risk, and controls, the AI Architect certification program builds on a legacy of high-quality training recognized worldwide. This certification provides you with an industry-respected credential that aligns with engineering standards, the EU AI Act, and global responsible technology frameworks, reinforcing your standing as a leader in AI governance and compliance.

## Testimonials

"*Congrats again on the training, really well received, i heard positive feedback across the board. My favorite in the material was the 100 risk scenarios you will run into. I printed that out and pinned on my wall. Great job*" - Ugur Ozer, AI Governance Lead, RBC Capital Markets, Canada

"*It was an intense three days of learning, capped off by three exams. The program provided a great opportunity to connect with professionals from around the world and dive deep into real-world AI governance, including some healthy debates on real-life AI use cases. The best part of the course was how practical it all was. It wasn't just theory, they provided real tools and frameworks that I can adapt and bring back to my work immediately*." - Ciara O'Buachalla, Co-founder & Chief Legal Officer @DonnaAI, Ireland

"*The workshop equipped me with invaluable tools, spreadsheets, guides, checklists, concepts, and even AI models ready to use in Python and test various scenarios and use cases, that will help me navigate the complexities of AI, from understanding challenges and requirements to ensuring compliance frameworks.*" - Aneta Klosek, Int. Dip, Director Global Trade, Germany

## World-Class Instructors

Our AI Architect certification program is led by distinguished experts with deep industry and academic experience, ensuring participants receive cutting-edge insights and mentorship from recognized leaders in AI, risk management, compliance, and corporate governance.

Program Director: Prof. Kersi F. Porbunderwalla

Prof. Porbunderwalla brings decades of expertise in governance, risk, and compliance rom both the corporate and academic worlds. As a former professor at IE Law School, Copenhagen Business School, and a lecturer at Georgetown University, Cass Business School, and Fordham University, he has influenced GRC education across institutions worldwide. He is the President and CEO of Copenhagen Compliance, The Information Security Institute, The Corporate Governance Institute, The EUGDPR, and The E-Compliance Academy. His work has led to the establishment of global GRC networks, offering certifications and training across four continents. Prof. Porbunderwalla's extensive experience in developing AI-driven solutions, GRC applications, and frameworks, particularly in specialized areas like cybersecurity and data privacy, ensures that program participants gain insights grounded in real-world applicability and regulatory standards.

Lead Trainer: Prof. Hernan Huwyler

Prof. Hernan Huwyler, MBA, CPA, CAIO, CCAR-P, trains teams on enterprise AI architecture using lessons from actual corporate deployments. He managed AI governance and risk control inside global technology and consulting firms. That background shapes his teaching approach. He focuses on what happens when models meet tight budgets, strict regulators, and real customer data. As an Academic Director at IE Business School and IE Law School, and lead instructor for the Chief AI Officer certification at Copenhagen Compliance, he works with cross-functional groups every week. Auditors demand evidence. Engineers require clear technical specifications. Executives care about cost and accountability. He structures his training programs to answer all three requirements directly. His training approach skips marketing hype to focus on concrete engineering

infrastructure, unit economics, and deterministic control boundaries. He holds the Anthropic Claude Certified Architect credential at the top professional level and teaches vendor-agnostic system design patterns that scale securely in production.

## Program Description

**Day 1: Architectural foundations, integration primitives and workflow orchestration**

Day 1 Learning objective Understand the foundational system design layers, integration protocols, enterprise primitives, and workflow orchestration patterns necessary to build secure, scalable, and business-aligned AI platforms.

09:00 AM – 10:15 AM | Executive welcome and module 1: The AI architect role, work decomposition and system foundations

Module and focus area: AI development planning and system boundaries.

Curriculum topics:

- Morning networking and opening remarks: Establishing systemic design paradigms and moving from probabilistic experimentation to deterministic enterprise reliability.
- The enterprise AI architect role: Translating ambiguous strategic imperatives into audit-ready architectural blueprints.
- Work decomposition framework: Dividing enterprise workflows strictly into probabilistic AI-assisted tasks, deterministic software-controlled tasks, and human-managed tasks. Deciding what the AI solution does, what plain software does, and what a person decides
- Avoiding the architecture trap: Why embedding deterministic business logic inside probabilistic model prompts leads to silent production failures.
- Requirements elicitation methodology: Uncovering unstated business, compliance, and operational constraints through structured stakeholder discovery (Case study: Logistics and claims automation).

10:15 AM – 10:30 AM | Morning coffee and executive networking break

10:30 AM – 12:00 PM | Module 2: Architectural isolation and the enterprise integration primitive stack

Module and focus area: System topology and integration engineering.

Curriculum topics:

- Three-tier architectural isolation: Separating user entry points, build-time developer interfaces, and cloud delivery routes (Amazon Web Services Bedrock, Azure AI Foundry, Google Cloud Vertex AI) to isolate risk and enforce compliance.

- The 7 enterprise primitives: Deep dive into tools, model context protocol, subagents, hooks, skills, agent teams, and dynamic workflows.
- Integration protocol stack: Programmatic mechanics comparing direct hypertext transfer protocol, application programming interfaces, language software development kits, agent orchestration SDKs, and model context protocol (MCP) server architectures (Case study: Commercial lease and contract review automation).

12:00 PM – 12:30 PM | Executive luncheon.

12:30 PM – 02:00 PM | Module 3: The autonomy spectrum and workflow orchestration sub-patterns

Module and focus area: Runtime execution and agentic design.

Curriculum topics:

- The autonomy spectrum: Evaluating operational trade-offs across augmented calls (zero autonomy), workflows (software-controlled pipelines), and autonomous agents (open-ended runtime loops).
- Workflow sub-patterns: Implementing prompt chaining (sequential), dynamic routing, parallelization (fan-out/fan-in), and evaluator-optimizer (generator-verifier) loops.
- Agentic state machine mechanics: Managing the four-phase runtime loop (plan, act, observe, stop) and establishing explicit termination criteria to prevent infinite execution traps.

02:00 PM – 02:15 PM | Afternoon refreshment break

02:15 PM – 04:00 PM | Module 4: Context engineering, memory architecture and system state

Module and focus area: Context state and memory management.

Curriculum topics:

- Working memory dynamics: Managing context window limitations, information decay across long transcripts, and attention-cliff phenomena.
- System context standardization: Designing persistent, versioned system instruction files to govern agent behavior across multi-turn sessions.
- Context lifecycle transitions: Balancing active working state, intermediate tool payload history, and automated background compaction.

04:00 PM | Day 1 Adjournment and executive reception

**Day 2: Economics engineering, governance, evaluation and executive defense**

Day 2 Learning objective Understand unit economics, multi-layer security guardrails, automated evaluation ladders, and architecture decision standards to deploy audit-ready, cost-effective enterprise AI platforms.

09:00 AM – 10:15 AM | Leadership briefing and Module 5: Economics engineering, prompt optimization and unit economics

Module and focus area: Cost optimization and token mechanics.

Curriculum topics:

- Executive recap and Day 2 plan: Aligning operational cost structures with enterprise risk governance.
- Context loading strategies: Technical trade-offs between monolithic context, progressive loading, just-in-time retrieval-augmented generation retrieval, and context compaction.
- Prompt caching engineering: Applying the static before dynamic structural rule, configuring time-to-live (TTL) parameters, and modeling cache setup surcharges against request volume.
- Unit economics and total cost of ownership modeling: Calculating direct infrastructure cost per completed business transaction, managing token throughput, and leveraging dynamic reasoning toggles (Case study: The unmanaged overcosts).

10:15 AM – 10:30 AM | Morning coffee and executive networking break

10:30 AM – 12:00 PM | Module 6: Enterprise governance, shared responsibility and threat control

Module and focus area: Security, compliance and human review architecture

Curriculum topics:

- Control layer ownership: Distinguishing provider trained behavior from operator system prompts and architect deterministic hooks.
- AI security threat mitigations: Defending against direct/indirect prompt injection, privilege escalation, data exfiltration, and skill supply chain vulnerabilities.
- Fail-closed guardrail design: Implementing mandatory automated validation gates and fail-closed isolation controls.
- Stakes-based human review routing: Designing human-in-the-loop checkpoints based on financial consequence, legal liability, and action reversibility.

12:00 PM – 12:30 PM | Executive luncheon

12:30 PM – 02:00 PM | Module 7: Model tiering, acceptance evals and the quality grading ladder

Module and focus area: Quality assurance and model evaluation

Curriculum topics:

- Model tiering strategy: Matching task complexity to flagship, mid-tier, or lightweight models to optimize performance against cost.
- The evaluation grading ladder: Implementing multi-tier test pipelines combining deterministic schema checkers, automated model-as-a-judge audits, and expert human reviews.
- Golden evaluation sets: Constructing representative, unbiased test suites containing edge cases, adversarial inputs, and real-world category proportions.
- Production observability: Tracking AI-specific failure taxonomies (hallucinations, schema breaches, unexpected refusals) via session correlation identifiers.

02:00 PM – 02:15 PM | Afternoon refreshment break

02:15 PM – 04:00 PM | Module 8: Architecture decision records, system deliverables and board-level defense

Module and focus area: Institutional memory and stakeholder alignment

Curriculum topics:

- Architecture decision records: Documenting design choices, rejected alternatives, and trade-off rationales to preserve institutional memory.
- The 8 standard system deliverables: Building the complete enterprise documentation package (AI system charter, architecture blueprint, architecture decision records, data documentation, model record, risk assessment, evaluation package, production monitoring contract).
- Stakeholder defense and objection resolution: Translating technical architectural decisions into business risk, cost, and compliance terms for executive leadership and audit boards. Best practices.
- Closing keynote

04:00 PM | Day 2 Adjournment

## Admission Process

To secure your place in the AI Architect certification program, we invite you to start by reaching out or submitting your application. You can request more information, including upcoming program dates, by contacting our team at eCompliance Academy. Simply email jeanette@e-compliance.academy for personalized assistance and guidance from a program advisor.

Upon receiving your application, our program coordinator will contact you by email to confirm receipt and guide you through the next steps. You'll be provided with a comprehensive overview of upcoming program dates and details on how to access your preparatory study materials. Payment details will be shared at this time as well, allowing you to secure your spot by fulfilling the program's financial requirements.

Please note that payment of program fees is required before the course begins and prior to receiving your pre-reading and study materials. This step is essential to confirm your participation and access the materials that will set the foundation for your learning.

We look forward to helping you advance toward an impactful career in AI leadership and management.