Deloitte.

Together makes progress

A Point of View on Security Architecture practices, their limitations, and adaptation to the AI era

Thoughts on the security-by-design process and its future evolution

Introduction

Cybersecurity that is discussed and designed at the inception stage is more durable, more affordable, and produces a stronger, more sustainable security posture. AI has made the timing critical and, at the same time, offers a real opportunity to formalize security architecture or uplift existing security architecture practices.

This Point of View focuses on one practice within the broader discipline of security architecture: security-by-design, sometimes referred to as "shifting security left." It is written for organizations at either end of the maturity range; those not yet applying it, and those running a version of it. The aim is practical: to help them uplift the practice, realize more value from it, and leverage the rise of AI to optimize and facilitate it while effectively managing the risks that the same technology introduces.

Section 1

Security architecture is a business decision

In an era where the cyber and information security mandate is no longer limited to a specific department or portfolio, the most consequential security decisions are now made before an information system is built, acquired, or subscribed to under an as-a-Service model. Cybersecurity mandate is shaped directly by business case reviews, funding approvals, procurement, contracting, third-party risk management, compliance, and at times law enforcement; activities that sit with departments well outside the traditional boundaries of IT and Cybersecurity. This holds true across corporate IT and Operational Technology alike, even where the specifics differ between the two. Security decisions, consciously or unconsciously, are made very early, often in the boardroom, and largely by people who sit in neither Cybersecurity nor IT.

Yet the practice meant to govern those decisions has not kept pace. Security architecture remains a loosely defined term, and is still approached, in most organizations, as a technical step the follows the finalization of the initial design and its functional and non-functional requirements. The security-by-design or "shift-left" movement is real and valuable, but it sits largely on the vendor's side of the responsibility model: it asks the maker to build securely and, too often, stops there. The buyer's side i.e. how a system is deployed, integrated, configured, and consumed in a specific environment, is assumed to be handled elsewhere, when formally it is frequently owned by no one. This is one of the foundational challenges architects face today.

Two everyday situations show how the gap widens. A commercial off-the-shelf (COTS) product may be genuinely well-built, with the vendor able to evidence secure development and testing, yet how it is hosted, hardened, integrated with internal systems, and operated inside the client's environment is often formally owned by no one, falling into the space between the vendor's assurances and the client's delivery teams. The second myth is that adopting SaaS or serverless technologies transfers security ownership entirely to the provider: "It runs on their platform, so it is their responsibility." In reality, key elements such as identity, configuration, data, integrations, and the way the service is deployed and consumed remain the customer's responsibility to design, manage, and secure. While the provider is responsible for securing the underlying platform, the customer must secure how the service is used in practice. This is simply an application of the shared responsibility model, explored further in Section 7, Governance: Who Owns This? While the vendor-side blind spot may be the most challenging, it doesn't stand alone. It is compounded by a practice still centered on software development rather than the full technology stack, and by governance that has not caught up with who actually owns these decisions, both examined in the sections that follow (Sections 3 and 7). And even where security-by-design is applied, it tends to be a point-in-time exercise rather than continuous security architecture assurance, which, in the age of AI-enabled systems, is no longer optional. The next section turns to what this delay costs, before Section 3 sets out why today's practice, valuable as it is, does not yet reach where these decisions are made.

Section 2

The cost of late: Why retrofitting security fails

The later security is considered, the more it costs and the less effective it becomes.

This is not a matter of opinion. NIST's foundational work on the economics of software quality established that the cost of correcting a defect rises sharply the later it is found in the lifecycle; Analyses based on that data suggest that production-stage fixes can cost more than 10 times as much as addressing them at an early stage, a cost differential that can reasonably be expected to be similar or even greater for security-related defects. The exact multiplier varies by project and organization, but the direction is uncontested[1]. This applies to correcting security defects, or more precisely, bolting security on.

While mature change management and agile delivery practices can absorb planned, calculated change, the effort and time required to retrofit security after the fact is of a different order entirely. Accepting the status quo can no longer be overlooked when the consequences directly degrade an organization's risk posture.

Reactive security routinely fails to deliver its real purpose: securing the organization and enabling the business to meet its objectives. Three pressures make this failure more acute than it once was:

First, dependence on technology continues to deepen, with AI now embedded across most digital and organizational transformations.

Second, organizational boundaries keep expanding through mergers, acquisitions, the dynamic third-party landscape, and the evolving international trading system.

Third, cross-jurisdiction compliance obligations are tightening.

One point deserves emphasis, because it sets up everything that follows: embedding security in information systems design and acquisition covers the full technology stack, not only the application layer. Delay in doing so will, more often than not, derail transformation plans, reducing competitive edge, costing money, and causing reputational damage that is hard to recover from.

[1]NIST (2002), The Economic Impacts of Inadequate Infrastructure for Software Testing.

Section 3

Security-by-design today: Why it falls short

It is imperative to note that a number of organizations already practice some form of security-by-design, and this current-state analysis is framed accordingly. The practice is real: it is actively advocated by organizations, teams, and individuals, and it is found in many mature organizations with aligned business units and security-informed strategies. Even so, current security-by-design practices tend to share five limitations.

- **Vendor-side centric:** pioneering advocacy, such as CISA's Secure-by-Design Pledge, asks makers to design and build with security in mind, but says little about how buyers acquire, configure, and place security across IT and IS acquisition more broadly. Yet securely built software is only as secure as the production state it lands in; how it is hosted, integrated, and consumed in a specific context that only the acquiring organization understands. When responsibility is framed as the maker's alone, the buyer's decisive role in that context goes ungoverned.

- **Software-engineering centric:** on a related note, most effort in this domain concentrates on software development and DevOps. But while the application layer (all tiers, front and back end) is often where attackers begin, they reach their ultimate objectives by exploiting the extended attack surface; the hosting, integrations, and building blocks customized after acquisition, or inherited from the wider environment, that are rarely examined thoroughly when the system is first designed. Securing the code while leaving that wider surface unexamined addresses only the smallest part of the problem.

- **Runs in silos:** the processes that should support security-by-design; TPRM, functional requirements definition, contracting terms and conditions, compliance in its wider (not only cyber) sense, and assorted technical and non-technical approval gates are run by different teams with different criteria, priorities, and perspectives. In organizations of every scale, it is common to find these checks neither centrally managed nor governed by an overarching framework. Without that synergy, a decision made at one gate can quietly undermine or duplicate one made at another.

- **Point-in-time:** whether the trigger is an internal milestone (security-by-design at concept stage, or a design review just before go-live) or a TPRM checkpoint (vendor assessment at onboarding), a reasonable level of targeted continuous assurance is often absent. Without it applied proportionately to system criticality, among other factors beyond this paper's scope, the value of security-by-design rarely survives past the initial release and never extends across the system's operational life.

- **Compliance-driven, not architecture-driven:** compliance checks and audits are powerful deterrent controls and should not be underestimated. But framing security-by-design as a compliance-mandated activity reinforces the point-in-time problem and defeats the purpose of proactive security: remediation tied to an audit observation can wait until the next assessment cycle, by which time the system has typically been in production for some time.

No single one of these limitations is a failure on its own; each reflects a practice that exists because it adds value. Taken together, though, they share one shape; security-by-design today tends to stop at the gate. The result is a practice that secures the moment of build or purchase, but rarely the life of the system that follows.

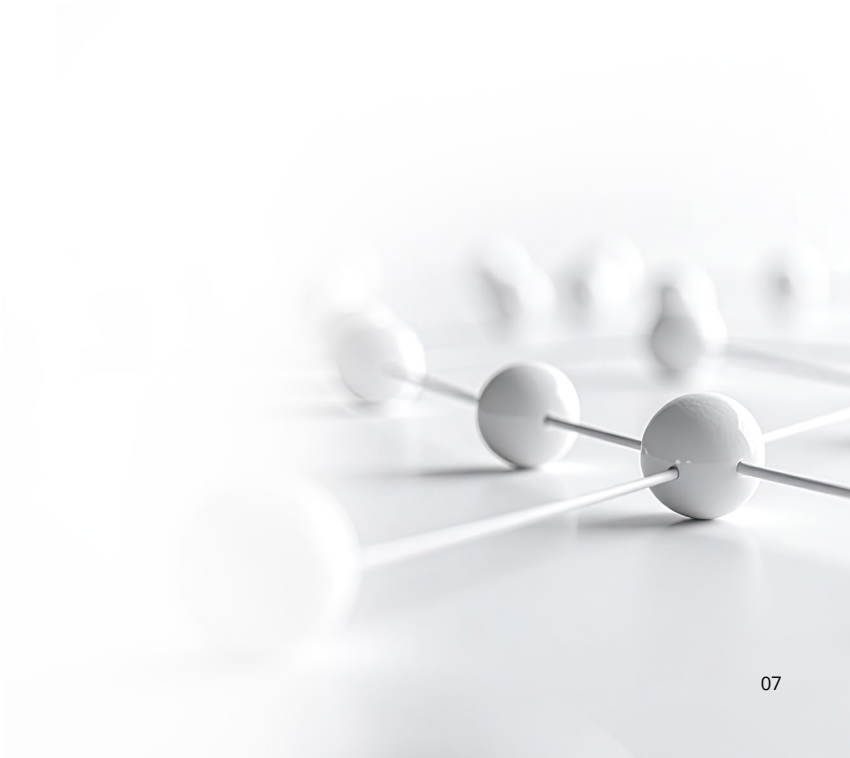

Section 4

The third-party and supply-chain reality

This is the first of two realities that make acting urgent rather than optional. It may sound obvious, but it sets up everything that follows: no organization operates autonomously. In the supply of IT goods and services, mandating source-side actions is not enough: the problem extends to the demand side, how a service is integrated, how it evolves, and how it eventually retires in the environment. That calls for assessing the full lifecycle of what is acquired, inside an integrated, organization-wide TPRM function rather than a cyber-only one, and for judging the security posture of the vendor, not just the product. A product built securely on day one can drift under the radar as its next release, or an as-a-Service offering, changes when the provider's own commitment slips. A point-in-time questionnaire will not catch that.

Relying on one anyway is, in effect, a decision to transfer the risk and keep a blame strategy in reserve for after a breach, rarely the wisest stance for a digitized, AI-embedded, data-driven organization. This is why continuous assurance belongs inside TPRM, and why security-by-design matters to it directly: continuous monitoring is expensive, and designing the requirement and the visibility in upfront is what makes assurance affordable at all. Today's mandates rarely provide for this. They seldom demand visibility into the vendor's own suppliers, fourth, fifth parties and beyond, yet a security-by-design tuned to TPRM would require exactly that knowledge, and would sometimes recommend not onboarding at all. Consider a provider that outsources its own IT and data management: the organization is not assessing one supplier, but an invisible chain of them.

Treated as a formal step that hands off cleanly into business-as-usual operations, supply-chain security pays back across asset management, monitoring, incident response, risk management, and continuous assurance.

Section 5

AI urgent reality

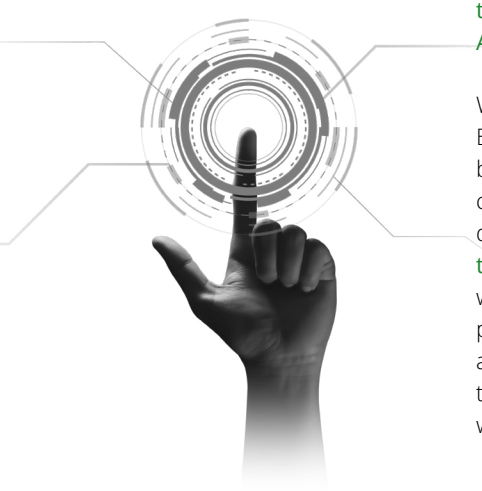

AI is the second reality confirming that action is not optional. While the world watches AI reshaping offensive and defensive security inside Security Operations Centers (SOC), a quieter shift gets far less attention: AI-enabled systems are already being supplied and used across organizations.

From narrow AI to frontier LLMs, these solutions have already begun to change the organizational threat landscape, creating a genuine need to address AI security architecture concerns at an early stage.

What makes AI different is that it does not hold still. By design it learns and adapts: models retrain and behavior drifts, and teams must continually correct confidently wrong outputs and control behavior that drifts toward pleasing the user over being right. A one-time design review doesn't govern a moving target; which makes this the sharpest case in the argument for pairing security-by-design with mandatory continuous assurance. It is also why control beats prohibition: better to adopt deliberately than to ban and scramble later, with capability now evolving in months, not years.

Security at the early stage of AI adoption mitigates risk:

- Architects who understand AI ask the right questions early: AI capability is routinely overstated in the sales pitch. An architect grounded in AI fundamentals can establish what a system actually is, narrow machine learning or a frontier model, GenAI, agentic, or multimodal, which changes the threat profile, the controls, and the resources required, and lets the capability grow with security built in rather than bolted on. Those same early questions also expose a quieter gap: the skills needed to run and assure the capability safely. Surfacing it at design rather than after go-live lets the organization decide in time whether to upskill its people or bring in a partner, and a partner is another third-party relationship to assure, not a way to hand the risk off.
- Design-in verifies that claimed AI controls actually work and tells SecOps what to watch: A vendor's claim that a control mitigates an AI risk is an assertion, not evidence; security-by-design is where it is verified before reliance. It also tells security operations what to monitor and where to focus threat intelligence and hunting, informed by organizational context the tooling alone will miss.

It also captures opportunity:

- Securely adopted AI improves the business processes it touches and, picking up the thread from Section 4, makes TPRM and continuous monitoring more effective and less costly. This puts third-party assurance at scale within reach of organizations previously blocked by cost.
- It lets security teams do slow, error-prone work faster; for instance, configuration hardening and continuous monitoring against an approved baseline, a high-value standardization task that time and overhead have long derailed.

AI does not replace security-by-design; it makes it more necessary, and at the same time offers a way to deliver it more easily. Either way, the practice can no longer be deferred.

Section 6

Standardization as a compounding advantage

Wherever it is possible, standardization is what makes the rest manageable. A standard, secure way of introducing change lets teams hold the stable parts of the estate to a rigid, continuously monitored baseline and spend their attention on the genuinely dynamic parts. The benefit is doubled for AI adoption, where operations teams are still mapping the solution stack: standardization shows them where security can be automated and where the fluid parts need watching. Security-by-design is what identifies where standardization is possible in the first place, often in infrastructure and supporting functions such as hosting and identity and access management.

The standardization call applies on both tracks. For internal initiatives, it strengthens security, simplifies operations, and lowers costs. For acquisitions, it sharpens TPRM and streamlines onboarding, monitoring, and offboarding; for example, a pre-approved control set attached to every vendor contract. It will not fit every case, but with planning, it fits far more often than not, and the balance improves over time: as more cases become standard patterns, the bespoke residue shrinks, and what remains is handled deliberately through governed, reactive change.

Standardization costs effort, but its benefit compounds. Once built, patterns are reused at scale, effectively free next to running full security-by-design from scratch each time, and the organization accrues a library that makes the practice faster, cheaper, and more consistent with every use. That consistency is itself a posture advantage: a standardized estate is easier to monitor, assure, and defend than a bespoke one. It is also what makes continuous security architecture assurance affordable at scale: you can't continuously assure infinite bespoke variety, but you can continuously assure a known set of approved patterns. And the relationship runs both ways: an AI-enabled security-by-design capability can flag standardization candidates for architects to validate, formalize, and test as contextualized patterns.

Section 7

Security architecture governance: Who owns the decision?

This Point of View has made the case for security-by-design's effectiveness, but the practice is, arguably, one of the most governance-dependent functions in cybersecurity. The silo problem named earlier is the root of it: many functions touch these decisions, each with its own priorities and timelines, and the change lands in sensitive domains where other people's KPIs are at stake. The last thing anyone wants is an assurance step that appears to cut across a committed delivery date.

That is exactly why the practice fails without clear ownership. When many functions touch the decision, but none owns the outcome, accountability falls through the cracks. Effective governance therefore requires top-down endorsement, the ultimate business decision-maker backing the need for the process, paired with a single named owner who holds the authority to make security-by-design and continuous assurance decisions stick at the point they are made. The answer is shared responsibility under single-point accountability. Ownership is not the same as doing everything: one accountable owner sits above a model in which each function carries a defined part.

- **Security teams (Architecture and SecOps):** sets cyber requirements and conduct cyber continuous assurance.
- **Enterprise architecture:** aligns those requirements to the overall design.
- **Procurement and Legal:** build them into sourcing, vendor selection, and contractual terms.
- **The business sponsor:** owns the risk decision.

Naming, and explicit allocation of these border lines, is what turns "everyone's concern" into someone's responsibility.

Section 8

Takeaways

Four recommendations follow from the argument so far, each addressing a challenge or reality raised earlier.

Re-allocate attention and time toward the design and acquisition phase: resist the urge to rush through design. Use this phase to work through the hard questions, security among them, but not only security. This is also where to make the foundational investment: an enterprise-wide security-by-design process, the standardization that makes it repeatable, and the AI-enabled tooling that helps run it consistently.

Mandate security requirements in procurement and third-party onboarding: including continuous assurance, not just point-in-time sign-off. The supply chain is where security posture is often inherited. Security requirements you define, continuously update, and regularly verify are what make that inheritance deliberate rather than accidental. These actions are typically carried out through a formal TPRM program.

Establish explicit ownership and governance for security-by-design and continuous assurance decisions: stand up and operate a security architecture working group; anchor it with strategy statements and security policies; and use clear responsibility matrices so the shared responsibilities named above are actually assigned.

Treat security as a design decision, not a bolt-on: this is the mindset that underpins the other three.

Section 9

The decision ahead

Security architecture is no longer a technical afterthought; it is a business design decision. Organizations that embed security requirements at the earliest stages of IT systems design, acquisition, and third-party onboarding, and that standardize the process effectively, are building a compounding security posture advantage. AI is raising the bar: the systems organizations are acquiring today are already AI-enabled, and the window to design security in is narrowing. Addressed proactively, that pressure becomes manageable, and the opportunity AI offers becomes capturable rather than lost.

This Point of View has explained the security-by-design value proposition in general terms and has shown it to be effective given the right executive buy-in and governance. As for where the practice should go from here: the organizations that get the most from it will extend security-by-design inward, applying to their own teams the same design discipline they demand of vendors. They will write it into the organization's own governance framework, not the cyber or technology governance most already have, but the enterprise-level governance where funding, procurement, and risk decisions are actually made, and keep raising awareness of why it matters. They will use AI deliberately, to relieve the very stress and cost that continuous assurance places on the effort rather than adding to it. For those with no security architecture practice, or only an ad-hoc one, this is the moment to establish it, and the views in this paper are a place to start.

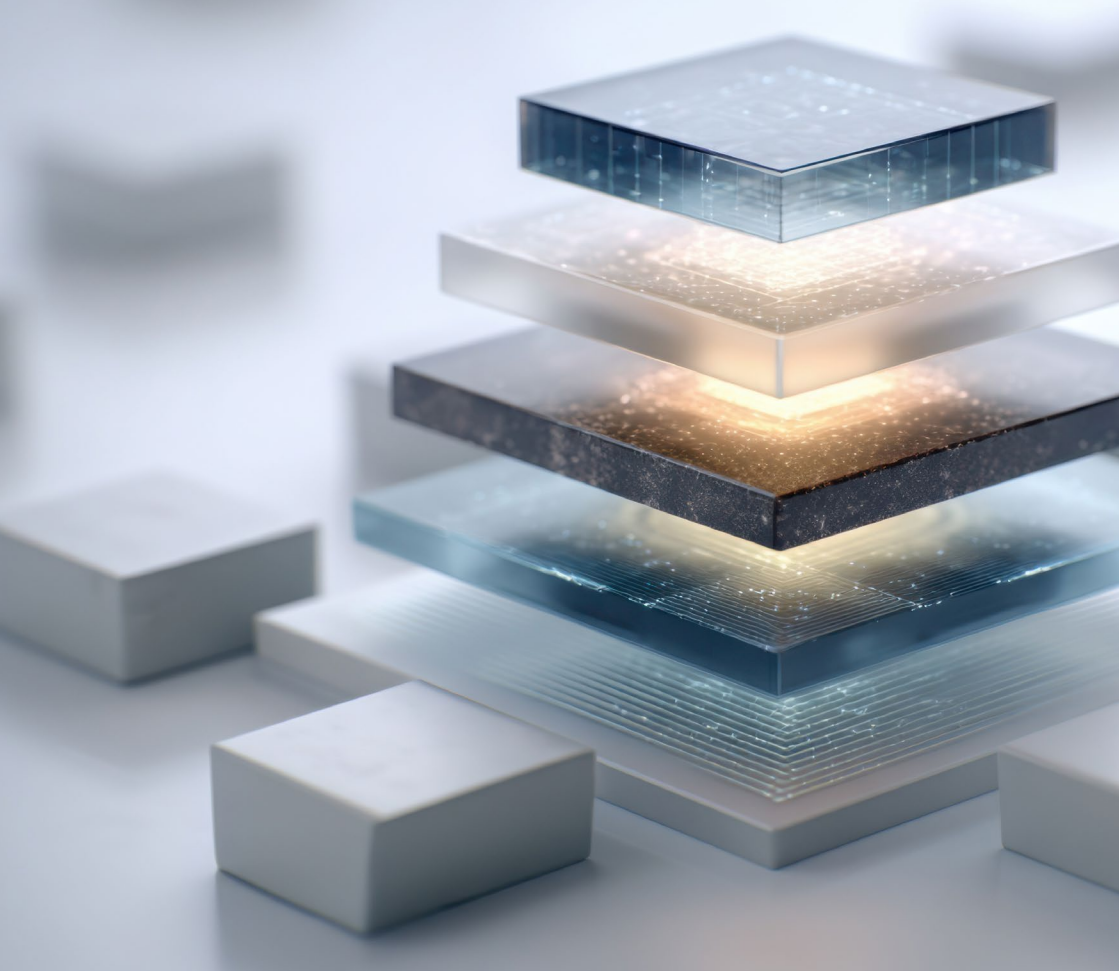

Authors:

Tamer Charife
Partner
Cyber Strategy and Transformation
Deloitte Middle East

Simsam Hijjawi
Director
Cyber Defense and Resilience
Deloitte Middle East

Deloitte.

Deloitte & Touche (M.E.) hereby authorizes you to view the information provided in this publication, subject to the following conditions: This publication has been written in general terms and therefore cannot be relied on to cover specific situations; application of the principles set out will depend upon the particular circumstances involved and we recommend that you obtain professional advice before acting or refraining from acting on any of the contents of this publication.
Deloitte & Touche (M.E.) (DME) is an affiliated sublicensed partnership of Deloitte NSE LLP with no legal ownership to DTTL. Deloitte North South Europe LLP (NSE) is a licensed member firm of Deloitte Touche Tohmatsu Limited.
Deloitte refers to one or more of DTTL, its global network of member firms, and their related entities. DTTL (also referred to as "Deloitte Global") and each of its member firms are legally separate and independent entities. DTTL, NSE and DME do not provide services to clients. Please see www.deloitte.com/about to learn more.

© 2026 Deloitte & Touche (M.E.) All rights reserved.

MAKING AN IMPACT THAT MATTERS since 1845