

# AI Copilots for PMI: Accelerating Workstreams with Evidence, Permissions and Human Sign-Off

*A Controlled Operating Framework for AI-Assisted Post-Merger Integration*

**Authored by Chennakeshav Adya, Independent Researcher**

September 2026

---

## Abstract

*Post-merger integration offices process large volumes of contracts, policies, process maps, issue logs, financial schedules, system inventories, communications and decisions under severe time pressure. AI copilots can improve retrieval, comparison, drafting, classification and issue triage, while also introducing risks of fabricated content, incomplete evidence, excessive access, confidential-information leakage, automation bias, uncontrolled model change and weak accountability. This paper develops a controlled operating framework for AI-assisted post-merger integration. It classifies use cases by consequence, establishes lawful and permissioned information boundaries, preserves source provenance, governs retrieval, prompts, models and tools, requires structured outputs, defines human review and sign-off, validates performance, monitors drift and incidents, and measures productivity without sacrificing decision quality. It addresses clean teams, privilege, personal data, cyber security, third-party providers, records retention, multilingual work, financial calculations, employee and customer communications, and board reporting. Five figures and five tables present the control chain, use-case risk map, evidence architecture, sign-off gate and certification structure. Eight frequently asked questions and forty primary or authoritative references support application. Numerical values are illustrative analytical scenarios. Transaction-specific conclusions require verified technical, legal, competition, privacy, cyber, employment, financial, operational and jurisdiction-specific evidence and advice.*

JEL Classification: G34, O33, M15, M42, D81

Keywords: AI copilots, post-merger integration, generative AI, evidence provenance, human oversight, permissions, clean teams, M&A

## 1. Establish accountable AI governance

The integration leadership should name integration, business, technology, security, privacy, legal, records and risk authority. The required output is an AI-copilot governance charter. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [1][2].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that experiments can become operational systems without accountable control. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

2. Inventory proposed use cases

The integration leadership should record task, users, data, output, downstream action, decision consequence and expected benefit. The required output is a controlled use-case register. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [1][3].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that a general tool approval can conceal materially different risks. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.


Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

Table 1. AI-copilot control chain

| Control layer | Evidence                   | Decision |
|---------------|----------------------------|----------|
| use case      | purpose and consequence    | classify |
| information   | sources and permissions    | permit   |
| output        | citations and validation   | review   |
| action        | human authority and record | approve  |

*Illustrative control structure; verified transaction evidence and specialist review govern.*

**Figure 1. Copilot control-readiness chain**



*Illustrative analytical scenario; verified transaction evidence should replace index values.*

### 3. Classify use cases by consequence

The integration leadership should rank assistance, recommendation, calculation, communication and decision influence. The required output is a consequence-tier matrix. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [1][4].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that low-friction drafting can influence high-consequence decisions. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

### 4. Define prohibited and restricted uses

The integration leadership should bar autonomous commitments, unapproved disclosure, privileged-data exposure and unsupported decision making. The required output is an acceptable-use boundary. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [2][5].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that speed incentives can bypass legal or fiduciary judgement. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **5. Protect pre-close competition boundaries**

The integration leadership should route competitively sensitive information through lawful clean teams and approved aggregation. The required output is an AI clean-team protocol. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [6][7].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that copilot retrieval can expose restricted information across deal teams. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **6. Map the information perimeter**

The integration leadership should identify repositories, data classes, owners, jurisdictions, retention and permitted purposes. The required output is a PMI information map. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [8][9].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that the model can access more information than the user may lawfully see. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **7. Enforce source-level permissions**

The integration leadership should inherit identity, role, matter, entity, geography and document restrictions at retrieval time. The required output is a permission enforcement design. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [10][11].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that a correct answer can still be an unauthorised disclosure. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **8. Minimise data exposure**

The integration leadership should limit prompts, retrieved context, outputs, logs and provider retention to the necessary information. The required output is a data-minimisation control. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [8][12].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that broad context windows can spread confidential or personal information. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **9. Protect legal privilege and confidentiality**

The integration leadership should segregate privileged matters, label restrictions and control derivative summaries. The required output is a privilege-preservation protocol. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [13][14].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that generated summaries can disclose protected analysis. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **10. Assess personal-data processing**

The integration leadership should define purpose, lawful basis, necessity, fairness, transparency, rights and human involvement. The required output is an AI data-protection assessment. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [8][15].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that integration analytics can repurpose employee or customer data unfairly. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## 11. Select models and providers

The integration leadership should evaluate capability, security, data use, residency, sub-processors, change control, exit and resilience. The required output is a provider assurance file. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [2][16].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that convenient tools can introduce hidden supply-chain exposure. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## 12. Document the system architecture

The integration leadership should map user, interface, retrieval, model, tools, storage, logging, controls and downstream systems. The required output is an auditable copilot architecture. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [2][17].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that unmapped components can evade testing and incident response. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

### **13. Govern prompts and instructions**

The integration leadership should version system prompts, templates, constraints, examples, prohibited actions and review dates. The required output is a prompt-control library. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [2][18].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that uncontrolled prompt edits can change the effective system. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

### **14. Ground outputs in retrieved evidence**

The integration leadership should require source citations, passage links, timestamps, permission context and uncertainty. The required output is a provenance-aware answer schema. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [1][19].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that fluent output can conceal missing or fabricated support. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.


Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

**Table 2. Evidence architecture**

| Element   | Minimum record         | Control   |
|-----------|------------------------|-----------|
| source    | document and version   | authorise |
| retrieval | passage and permission | trace     |
| output    | claim and uncertainty  | validate  |
| decision  | reviewer and action    | retain    |

*Illustrative control structure; verified transaction evidence and specialist review govern.*

**Figure 2. Evidence traceability**



*Illustrative analytical scenario; verified transaction evidence should replace index values.*

## 15. Control document ingestion

The integration leadership should verify source authenticity, version, completeness, malware status, classification and access. The required output is an ingestion-control gate. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [20][21].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that corrupt, stale or hostile documents can contaminate outputs. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **16. Use structured output schemas**

The integration leadership should define required fields, evidence, confidence, exceptions, owner and next action. The required output is a machine- and human-readable result. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [1][22].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that free-form answers can omit critical integration facts. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **17. Constrain tool use and agency**

The integration leadership should whitelist retrieval, calculation, drafting and workflow actions with bounded permissions. The required output is a tool-authority matrix. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [2][23].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that an assistant can mutate systems or transmit information beyond intent. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## 18. Separate drafting from approval

The integration leadership should ensure generated work remains visibly pending until an authorised person reviews and releases it. The required output is a human sign-off gate. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [4][24].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that users can mistake generated content for an approved decision. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.


Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

**Table 3. Human sign-off gate**

| Review dimension | Question                 | Evidence        |
|------------------|--------------------------|-----------------|
| facts            | supported and current?   | citations       |
| calculation      | reproducible?            | workings        |
| authority        | reviewer qualified?      | approval        |
| consequence      | downstream impact known? | decision record |

*Illustrative control structure; verified transaction evidence and specialist review govern.*

**Figure 3. Human sign-off strength**



*Illustrative analytical scenario; verified transaction evidence should replace index values.*

## 19. Design substantive human review

The integration leadership should require verification of facts, calculations, completeness, uncertainty, conflicts and consequences. The required output is a reviewer checklist. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [15][25].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that human involvement can become a ceremonial click. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## 20. Match reviewers to expertise

The integration leadership should route finance, legal, cyber, tax, HR, operational and customer outputs to qualified approvers. The required output is an expertise-routing map. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [1][26].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that a reviewer can approve outside their competence. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **21. Validate retrieval quality**

The integration leadership should test recall, precision, permission enforcement, version selection and citation correctness. The required output is a retrieval evaluation pack. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [1][19].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that the system can miss decisive evidence while appearing complete. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **22. Test factual and calculation accuracy**

The integration leadership should use representative, adversarial and boundary cases with independently verified answers. The required output is an output-validation suite. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [1][27].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that plausible errors can propagate into plans and forecasts. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **23. Test bias and differentiated outcomes**

The integration leadership should compare error, escalation and decision effects across relevant groups and contexts. The required output is an outcome-testing record. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [3][15].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that automation can amplify uneven data or historic practice. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **24. Test prompt injection and data exfiltration**

The integration leadership should challenge documents, links, instructions, tools and cross-context access. The required output is an AI security test pack. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [2][28].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that hostile content can redirect the copilot or reveal protected information. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **25. Control model and configuration change**

The integration leadership should version models, retrieval indexes, prompts, policies, tools and evaluation thresholds. The required output is a controlled release record. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [2][29].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that provider or configuration change can invalidate prior assurance. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **26. Create workstream-specific playbooks**

The integration leadership should define permitted tasks, evidence, reviewer, turnaround, escalation and retained records. The required output is a bounded copilot playbook. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [1][22].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that generic deployment can ignore functional risk. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **27. Accelerate document comparison**

The integration leadership should extract and compare terms, policies, processes and controls with source-level exceptions. The required output is a verified difference register. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [13][19].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that summaries can flatten legally or operationally material differences. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## **28. Support issue and dependency management**

The integration leadership should classify issues, connect evidence, identify owners and suggest escalation without closing decisions. The required output is an evidence-linked issue log. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [22][30].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that automated status can obscure unresolved judgement. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## 29. Support financial analysis

The integration leadership should use controlled calculations, source cells, assumptions, reconciliations and finance approval. The required output is a calculation evidence pack. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [27][31].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that generated arithmetic or mappings can distort value and liquidity. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## 30. Control employee and customer communications

The integration leadership should verify facts, tone, audience, timing, legal obligations and accountable release. The required output is a communication approval record. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [24][32].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that generated messaging can create commitments or confusion. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

### **31. Govern multilingual work**

The integration leadership should test terminology, legal meaning, numerical fidelity and culturally sensitive communication. The required output is a language-quality control. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [1][33].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that translation fluency can conceal material semantic error. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

### **32. Log material interactions**

The integration leadership should retain task, user, sources, permissions, model, prompt, output, review, action and exception. The required output is an auditable interaction record. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [2][34].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that important decisions can lack reproducible evidence. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

### 33. Monitor production performance

The integration leadership should track accuracy, citation quality, permission failures, overrides, incidents, latency and usage. The required output is an AI control dashboard. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [1][2].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that adoption metrics can conceal declining reliability. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.


Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

**Table 4. Production monitoring**

| Indicator         | Signal                 | Response    |
|-------------------|------------------------|-------------|
| citation failure  | unsupported claim      | block       |
| permission breach | unauthorised retrieval | contain     |
| override rate     | reviewer rejection     | investigate |
| quality drift     | evaluation decline     | revalidate  |

*Illustrative control structure; verified transaction evidence and specialist review govern.*

**Figure 4. Production-control coverage**



*Illustrative analytical scenario; verified transaction evidence should replace index values.*

### **34. Manage incidents and containment**

The integration leadership should define reporting, isolation, access revocation, evidence preservation, notification and recovery. The required output is an AI incident plan. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [2][35].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that a model or data event can spread across integration workstreams. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

### **35. Provide challenge and appeal**

The integration leadership should allow users and affected teams to question outputs, correct sources and escalate decisions. The required output is an output-challenge path. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [15][25].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that automation bias can suppress contrary evidence. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

### **36. Measure productivity with quality**

The integration leadership should compare cycle time, rework, error, reviewer effort, decision delay and realised outcome. The required output is a balanced value scorecard. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [1][36].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that token or usage volume can be mistaken for integration value. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

### **37. Report material use to leadership**

The integration leadership should present use cases, exposure, performance, incidents, changes, benefits and residual risk. The required output is a board AI-assurance view. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [3][37].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that leadership can lack visibility over consequential AI dependence. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

### **38. Retire unsafe or obsolete uses**

The integration leadership should withdraw access, preserve records, migrate work, verify deletion and communicate residual actions. The required output is a controlled retirement record. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [2][38].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that abandoned pilots can retain data or influence. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

### **39. Issue the copilot-use certificate**

The integration leadership should reconcile task, authority, sources, permissions, versions, validation, review, decision and incidents. The required output is an auditable copilot-use certificate. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [1][2].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that operational reliance can outlive its evidence. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.


Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

**Table 5. Copilot-use certificate**

| Conclusion      | Evidence                 | Acceptance |
|-----------------|--------------------------|------------|
| bounded purpose | use case and owner       | approved   |
| controlled data | sources and permissions  | verified   |
| reliable output | tests and review         | accepted   |
| residual risk   | incidents and exceptions | owned      |

*Illustrative control structure; verified transaction evidence and specialist review govern.*

**Figure 5. Copilot certificate readiness**



*Illustrative analytical scenario; verified transaction evidence should replace index values.*

## 40. Institutionalise safe AI-assisted integration

The integration leadership should retain evaluations, incidents, reviewer feedback, benefits and reusable controls. The required output is an enduring AI integration control cycle. Record the business purpose, task consequence, information perimeter, accountable owner, technical configuration, evidence requirement, human authority, control implications and review date [1][39].

Test the copilot across both businesses and the complete information route. Verify identity, permissions, source version, retrieval quality, prompt and model configuration, output schema, citations, calculations, uncertainty, reviewer competence and downstream action. Apply documented human judgement to verified evidence.

The principal risk is that each transaction can repeat preventable deployment failures. Compare prevention, detection, containment, recovery and challenge. Require the system to expose its sources and limitations, abstain when evidence is insufficient and keep operational decisions with authorised people.

Retain use-case approvals, data maps, provider assurance, prompt and model versions, evaluations, retrieved sources, generated outputs, reviewer actions, downstream decisions, logs, incidents, changes and exceptions. Revalidate when data, access, models, prompts, tools, users, law or transaction conditions change.

## References

- [1] National Institute of Standards and Technology, AI Risk Management Framework, <https://www.nist.gov/itl/ai-risk-management-framework>
- [2] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework Generative AI Profile, <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- [3] European Commission, AI Act, <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- [4] European Commission, Transparency Obligations under Article 50 of the AI Act, <https://digital-strategy.ec.europa.eu/en/faqs/transparency-obligations-under-article-50-ai-act>
- [5] European Commission, Navigating the AI Act, <https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act>
- [6] Federal Trade Commission, Avoiding Antitrust Pitfalls during Pre-Merger Negotiations and Due Diligence, <https://www.ftc.gov/enforcement/competition-matters/2018/03/avoiding-antitrust-pitfalls-during-pre-merger-negotiations-due-diligence>
- [7] UK Competition and Markets Authority, Interim Measures in Merger Investigations, <https://www.gov.uk/government/publications/interim-measures-and-derogations-guidance-and-templates>
- [8] UK Information Commissioner's Office, Guidance on AI and Data Protection, <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/>
- [9] UK Government, Government Data Quality Framework, <https://www.gov.uk/government/publications/the-government-data-quality-framework/the-government-data-quality-framework>
- [10] National Cyber Security Centre, Zero Trust Architecture Design Principles, <https://www.ncsc.gov.uk/collection/zero-trust-architecture>
- [11] National Institute of Standards and Technology, SP 800-207 Zero Trust Architecture, <https://csrc.nist.gov/pubs/sp/800/207/final>
- [12] European Union, General Data Protection Regulation, <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- [13] UK Government, Legal Professional Privilege Guidance, <https://www.gov.uk/government/publications/legal-professional-privilege>
- [14] UK Information Commissioner's Office, Data Sharing Code of Practice, <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/data-sharing-a-code-of-practice/>
- [15] UK Information Commissioner's Office, Automated Decision-Making and Profiling, <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/individual-rights/automated-decision-making-and-profiling/>
- [16] National Cyber Security Centre, Supply Chain Security Guidance, <https://www.ncsc.gov.uk/collection/supply-chain-security>
- [17] National Cyber Security Centre, Guidelines for Secure AI System Development, <https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development>
- [18] National Institute of Standards and Technology, AI RMF Playbook, <https://airc.nist.gov/airmf-resources/playbook/>
- [19] National Institute of Standards and Technology, Information Quality Act Resources, <https://www.nist.gov/director/information-quality-standards>
- [20] National Cyber Security Centre, Secure Development and Deployment Guidance, <https://www.ncsc.gov.uk/collection/developers-collection>

- [21] Cybersecurity and Infrastructure Security Agency, Secure by Design, <https://www.cisa.gov/securebydesign>
- [22] UK Government Service Manual, Agile Delivery, <https://www.gov.uk/service-manual/agile-delivery>
- [23] National Institute of Standards and Technology, Secure Software Development Framework, <https://csrc.nist.gov/projects/ssdf>
- [24] European Commission, Commission Starts Enforcing AI Act Rules and Transparency Requirements, <https://digital-strategy.ec.europa.eu/en/news/commission-starts-enforcing-ai-act-rules-and-new-transparency-requirements-2-august>
- [25] OECD, OECD AI Principles, <https://oecd.ai/en/ai-principles>
- [26] International Organization for Standardization, ISO IEC 42001 AI Management Systems, <https://www.iso.org/standard/81230.html>
- [27] UK Statistics Authority, Code of Practice for Statistics, <https://code.statisticsauthority.gov.uk/>
- [28] National Cyber Security Centre, Prompt Injection Attacks and Mitigations, <https://www.ncsc.gov.uk/blog-post/prompt-injection-attacks-and-mitigations-1>
- [29] National Cyber Security Centre, Machine Learning Principles, <https://www.ncsc.gov.uk/collection/machine-learning-principles>
- [30] UK Government Infrastructure and Projects Authority, Project Delivery Functional Standard, <https://www.gov.uk/government/publications/project-delivery-functional-standard>
- [31] IFRS Foundation, Conceptual Framework for Financial Reporting, <https://www.ifrs.org/issued-standards/list-of-standards/conceptual-framework/>
- [32] UK Government Service Manual, Writing for GOV.UK, <https://www.gov.uk/guidance/content-design/writing-for-gov-uk>
- [33] International Organization for Standardization, ISO 17100 Translation Services, <https://www.iso.org/standard/59149.html>
- [34] National Cyber Security Centre, Logging Made Easy, <https://www.ncsc.gov.uk/information/logging-made-easy>
- [35] National Cyber Security Centre, Incident Management, <https://www.ncsc.gov.uk/collection/incident-management>
- [36] UK Government, The Green Book, <https://www.gov.uk/government/publications/the-green-book-appraisal-and-evaluation-in-central-government>
- [37] Financial Reporting Council, UK Corporate Governance Code 2024, <https://www.frc.org.uk/library/standards-codes-policy/corporate-governance/uk-corporate-governance-code/>
- [38] UK Government Service Manual, Retiring Your Service, <https://www.gov.uk/service-manual/technology/retiring-your-service>
- [39] Committee of Sponsoring Organizations of the Treadway Commission, Internal Control Integrated Framework, <https://www.coso.org/guidance-on-ic>
- [40] European Commission, Guidelines on Obligations for General-Purpose AI Providers, <https://digital-strategy.ec.europa.eu/en/faqs/guidelines-obligations-general-purpose-ai-providers>

## About the Author

Authored by Chennakeshav Adya, Independent Researcher

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds; bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.