

FREE SUMMARY

# AE-OS

## Agent Controls Pack

Controls for AI agents that post journals and release payments in your finance systems.

Written for CFOs, controllers, internal audit and audit committees.

REFERENCED TO

|                                                              |                                                       |                                                         |                                              |
|--------------------------------------------------------------|-------------------------------------------------------|---------------------------------------------------------|----------------------------------------------|
| <b>COSO 2013</b><br>INTERNAL CONTROL<br>FRAMEWORK PRINCIPLES | <b>NIST AI RMF</b><br>AI RISK MANAGEMENT<br>FUNCTIONS | <b>ISO/IEC 42001</b><br>AI MANAGEMENT SYSTEM<br>CLAUSES | <b>EU AI Act</b><br>ARTICLES, WHERE RELEVANT |
|--------------------------------------------------------------|-------------------------------------------------------|---------------------------------------------------------|----------------------------------------------|

|                                                                                                                     |                                                                                                                                |                                                                                                                |
|---------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------|
| <b>FOR CFOS AND CONTROLLERS</b><br>Which controls must be in place before agents post journals or release payments. | <b>FOR INTERNAL AUDIT</b><br>Scope agent risks into the audit plan with suggested tests of design and operating effectiveness. | <b>FOR AUDIT COMMITTEES</b><br>Ask management the questions that show whether agent activity is under control. |
|---------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------|

01

WHY IT MATTERS

# AI agents and internal control over financial reporting

An AI agent can read an invoice, match it to a purchase order and propose the journal. In some setups it also posts the entry to SAP, Oracle, NetSuite or Dynamics and releases the payment. Auditors and audit committees will ask the same questions about those steps that they ask when a person takes them: who acted, under what authority, with whose approval, and what evidence shows it.

Internal control over financial reporting (ICFR) was designed around human users and conventional programs. An agent makes choices and enters transactions as a user does. Its behavior changes when someone edits its instructions or the provider releases a new model version, much as a program changes when new code is released. Existing access and change controls cover part of this risk, and the pack sets out the additional controls needed.

## What goes wrong when agents run under existing controls

|                                                                                                                                            |                                                                                                                                                               |
|--------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Attribution</b><br>If an agent posts under a shared or human account, your audit trail shows the wrong actor.                           | <b>Authority</b><br>If an agent can do everything its service account can do, your delegation of authority matrix is bypassed.                                |
| <b>Change</b><br>If someone edits an agent's instructions without approval, a key automated control can change without anyone signing off. | <b>Evidence</b><br>If the record of what the agent did can be edited, auditors will place less reliance on it, and testing becomes slower and more expensive. |

THE FULL PACK

The AE-OS Agent Controls Pack sets out controls for each of these risks. Every control has a risk statement, a control activity, suggested audit tests and the evidence to keep, in a format that fits an existing risk and control matrix.

## How agents compare with people and conventional programs

|                       | PEOPLE                         | CONVENTIONAL PROGRAMS                   | AI AGENTS                                                        |
|-----------------------|--------------------------------|-----------------------------------------|------------------------------------------------------------------|
| Make choices          | Yes                            | No, they follow fixed logic             | Yes, within the scope they are given                             |
| Enter transactions    | Yes                            | Yes                                     | Yes                                                              |
| Behavior changes when | Training, turnover             | Code is released through change control | Instructions, tools or the model version change                  |
| Usual controls        | Access reviews, approvals, SoD | ITGCs over change and operations        | Both sets, extended to agent identities, instructions and models |

## 02 CONTROL DOMAINS

# Control domains in the pack

The pack groups its controls by domain. The examples below show the kind of control each domain contains.

## A Agent identity and access

- Give each agent its own non-human identity in your identity system and in every ERP it touches, so every action is attributable.
- Issue agents short-lived credentials from a secrets manager, kept out of prompts and logs.
- Have each agent act on behalf of a named person, limited to the lesser of its own permissions and that person's authority, and include agents in quarterly access reviews.

## B Change management for agents

- Treat prompts and instructions as program code: version control, peer review, testing against finance test cases and owner approval.
- Pin the exact model version in production and re-test before any upgrade, including upgrades announced by the provider.
- Require finance control owners to approve changes to limits and approval thresholds in the agent's policy rules.

## C Authority and application controls

- Enforce value limits per transaction, per vendor and per period in both the policy layer and the ERP.
- Extend segregation of duties to agents. An agent acting for a person counts as that person, and two agents may not combine incompatible steps.
- Route transactions above thresholds, with new counterparties or touching sensitive accounts to a human approver, and keep agents away from vendor master and bank detail changes.

## D Operations and monitoring

- Record every agent action with who, what, when, on whose behalf and under which policy decision, and block the action if logging fails.
- Make the log tamper-evident and have someone independent verify its integrity regularly.
- Keep a tested way to suspend any agent immediately, and reconcile agent logs to ERP postings daily.

## E Third-party model and tool providers

- Assess providers before onboarding for security, data use, model change notification and exit terms.
- Review SOC 1 or SOC 2 reports and map the complementary user entity controls to your own controls.
- Document and exercise a fallback for each agent that supports close or payment runs.

## F Data and privacy

- Give agents only the data fields their task needs.
- Set a retention schedule for prompts, outputs and logs that meets both audit and privacy requirements.
- Block unapproved AI tools and define what confidential data may be included in prompts.

## G Governance

- Maintain an inventory of every agent in or near finance processes, with a named finance owner and technical owner.
- Complete a risk assessment before deployment covering ICFR impact, fraud risk and regulatory classification.
- Evaluate each agent's accuracy and controls at least annually and report the results to the audit committee.

## H Reporting and close

- Park agent-prepared journals for approval by someone independent of the agent and the person it acts for.
- Require agent-prepared reconciliations to use validated system balances and a qualified reviewer's sign-off.
- Allow agents to prepare estimates only with approved methods, keep their inputs for re-performance, and have management review the assumptions.

**03 QUESTIONS TO ASK**

# Questions to ask before an agent posts to the ledger

Put these to the team proposing the agent. Clear answers show that the controls are designed, and testing then shows whether they operate effectively.

**01** Does this agent have its own identity in the ERP, separate from any person or shared service account?  
IDENTITY · DOMAIN A

**02** Which named person is the agent acting for, and does that person have authority for every transaction the agent can make?  
DELEGATION · DOMAIN A

**03** What exactly can the agent do, in which company codes and up to what amount, and where is that written down and approved?  
SCOPE AND LIMITS · DOMAINS A AND C

**04** Which transactions does the agent send to a person for approval, and can the agent complete them without that approval?  
ESCALATION · DOMAIN C

**05** Can the agent create a vendor, change bank details or approve its own work? In most finance processes each of these should sit with a person or a separately controlled process.  
SEGREGATION OF DUTIES · DOMAIN C

**06** Who can change the agent's instructions, model version or rules, and what testing and approval do they need first?  
CHANGE MANAGEMENT · DOMAIN B

**07** Is every agent action recorded, and can we prove the record has not been altered since it was written?  
LOGGING AND INTEGRITY · DOMAIN D

**08** How do we reconcile what the agent says it did with what is actually in the ledger?  
COMPLETENESS · DOMAIN D

**09** How quickly can we stop the agent, and when did we last test that?  
SUSPENSION · DOMAIN D

**10** What happens to close or the payment run if the AI provider is unavailable for a day?  
RESILIENCE · DOMAIN E

**IF AN ANSWER IS UNCLEAR**

Start the agent in a mode where it prepares work for a person to post, and move to posting only when the controls are in place and tested.

04

THE FULL PACK

# What the full pack contains

The full AE-OS Agent Controls Pack is an Excel workbook structured to be added to an existing risk and control matrix.

|                         |                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RISK AND CONTROL MATRIX | 45 controls, grouped by domain. Each includes the risk, control objective, control activity, type and nature, frequency, suggested key control flag, related assertion or ITGC area, suggested test of design, suggested test of operating effectiveness and evidence to retain.                                                                                |
| FRAMEWORK REFERENCES    | COSO 2013 principles, NIST AI RMF functions, ISO/IEC 42001 clauses (with Annex A controls left to map in your own Statement of Applicability) and EU AI Act articles where relevant. The EU AI Act references apply in full only to high-risk systems; many finance agents fall outside that category, and the references then serve as good-practice guidance. |
| EVIDENCE NOTES          | Where a tamper-evident, signed per-action log (such as AE-OS Receipts) provides audit evidence. Where such a log adds no evidence for a control, the pack says so.                                                                                                                                                                                              |
| MAPPING SUMMARY         | Counts by domain, control type, nature and framework, to support scoping and coverage discussions.                                                                                                                                                                                                                                                              |
| GLOSSARY                | Plain-language definitions of terms such as non-human identity, capability scoping, on-behalf-of delegation, inclusion proof and consistency proof.                                                                                                                                                                                                             |
| GUIDANCE                | How to scope the pack to your agents, tailor thresholds and frequencies, and plan testing.                                                                                                                                                                                                                                                                      |

## Extract from the risk and control matrix

| ID    | DOMAIN                 | CONTROL ACTIVITY                                                                                                                                                                                                                       | TYPE       | NATURE              | KEY | COSO       |
|-------|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|---------------------|-----|------------|
| AC-01 | A. Identity and access | Each agent is issued its own non-human identity in the identity provider and in each connected ERP. Shared accounts and use of human credentials by agents are prohibited and technically blocked where the platform allows.           | Preventive | Automated           | Y   | 10, 11     |
| AU-04 | C. Authority           | Escalation rules route defined cases (value above threshold, new counterparty, manual journal to sensitive accounts, low agent confidence) to a named human approver. The agent cannot complete the action until approval is recorded. | Preventive | IT-dependent manual | Y   | 10, 12     |
| OM-02 | D. Operations          | Log entries are signed and appended to a tamper-evident structure. Inclusion and consistency proofs are verified periodically by a party independent of agent operations.                                                              | Detective  | Automated           | Y   | 11, 16     |
| RC-01 | H. Close               | Agent-prepared journals are parked for approval before posting. A reviewer independent of the preparer and its principal reviews the entry, its support and the agent's explanation.                                                   | Preventive | IT-dependent manual | Y   | 10, 12, 13 |

Control activities shortened for this page. The full workbook adds risk statements, test steps, evidence and framework references for every control.

### OBTAINING THE FULL PACK

Licensed per organization for internal use, including sharing with your external auditors.

ae-os.org

### IMPORTANT NOTE

This summary and the full pack are general guidance and do not constitute legal, audit, accounting or regulatory advice. Each organization should tailor the controls to its own processes, systems, risks and obligations, and apply professional judgment. Confirm the regulatory classification of any AI system with legal counsel.