

Segregation of duties across the whole estate


One rule engine and one evidence model spanning your ERP, HR and SaaS applications — and the identities that aren't people. It runs alongside the identity platform you already own rather than replacing it.

- Cross-application segregation of duties
- People, service accounts and AI agents
- Preventive and detective
- Works with your existing IAM

THE PROBLEM

Segregation of duties does not stop at the application boundary

Each platform models access differently, and each one polices only itself. The combination that actually enables fraud is usually assembled from parts that are individually reasonable — and no single application can see it.



Three applications, three clean-looking access reviews, one complete fraud path. Evaluating it requires a rule that spans all three — and a model in which a Workday domain grant, a Fusion privilege and an EBS function are the same kind of object.

COVERAGE

Built for how each platform actually grants access

Rules evaluate the **effective permission** — after inheritance and nesting are resolved — and then normalise it. Normalisation is what makes a cross-application rule possible at all.

PLATFORM	ACCESS IS GRANTED BY	RESOLVED TO	DATA DIMENSION
Oracle EBS	Responsibility	Menu and form function	Operating unit
Oracle Fusion	Job and abstract role	Duty role, then privilege	Business unit, ledger
SAP	Single and composite role	Object, field and value	Company code, plant
Workday	Security group	Domain and business-process policy	Organisation
PeopleSoft	Role	Permission list	Business unit
Dynamics 365, NetSuite	Security role	Duty and privilege	Legal entity, subsidiary
SaaS and custom applications	Application role or profile	Permission, via configurable extract	As modelled

PLATFORM	ACCESS IS GRANTED BY	RESOLVED TO	DATA DIMENSION
Machine identities	API permission, agent tool grant	Callable operation	Environment

Once an EBS function, a Fusion privilege, a SAP value and a Workday domain grant are expressed as the same kind of object, one rule spans all of them — and one exception, one control and one audit record follow it.

CAPABILITIES

What the risk engine does

Fine-grained risk analysis Analyse at whichever grain the rule needs: the role as granted, the business function beneath it, or the individual permission and value. Coarse rules that keep the register readable and surgical rules that remove false positives run side by side in the same analysis.	Pre-built rule libraries Ships with a controls library of more than nine thousand rules spanning Financials, Human Capital, Procurement and Order-to-Cash, so a first meaningful risk picture arrives in days rather than quarters. Your own matrix loads directly from a spreadsheet, and any library exports for review or handover.
Mitigating control recommendations Controls are first-class records — preventive or detective, with a test frequency and a named owner — pre-linked to the rules they mitigate. When a violation surfaces, the recommended control set is already attached, so the reviewer assigns rather than researches. Assignments carry effective dates and lapse on schedule.	Segregation of duties and risk policy review A reviewer workqueue with ownership, bulk disposition, threaded discussion and a full evidence trail behind every finding: which identity, which role, which permission, which scope, which rule. Role-pair review shows which combinations can conflict before anyone is assigned to them.
Inheritance lineage Every finding opens to the complete chain from person to permission, and where a permission arrives by several routes at once, all of them are shown. Revoking a single role often does not clear a violation — lineage names the specific grants that do, so remediation is surgical instead of removing a role and breaking someone's job.	Audit-ready reporting Every analysis is a dated, durable record of its ruleset, its population and its results — reproducible months later for a walkthrough. Findings export with their evidence trail intact, rule libraries export as point-in-time snapshots, and the change history exports alongside them.
What-if analysis for role design Four modes of simulation: add or remove access for an existing identity; model a new hire from an empty baseline; test a single role for conflicts it carries on its own; and compare two roles to separate intrinsic risk from risk that appears only in combination. Role-conflict results are precomputed, so redesign stays interactive across thousands of roles.	Compensating controls and exceptions Exceptions are requested with a written justification, approved or rejected by a separate party, and bounded by a validity window with its own recertification date. Nothing is permanent by default — expiry runs on schedule and the violation returns to the queue for a fresh decision.

Continuous monitoring of high-risk access Two enforcement points, not one. Scheduled analysis re-scores the whole population against the current rules; a request-time check evaluates the conflict <i>before</i> access is granted, so provisioning is blocked or escalated rather than remediated later. New high and critical findings are pushed to your monitoring, ticketing and alerting systems as they arise.	Multiple rule libraries, versioned Run as many libraries as the business needs — external audit, internal audit, a regional variant, a revision under review — at the same time and against the same population. A library clones in a single step to create the next version, revisions stay in draft until promoted, and shared rules mean a correction authored once propagates everywhere it belongs.
Approvals on risk register changes Exceptions and compensating control assignments move through request, review and decision with a named requester, a named approver, a written justification and a bounded validity period. Rule and library revisions are authored as drafts and promoted deliberately, so nothing reaches the live register by accident.	Complete history of register changes Every change to the risk register is captured with the full before and after state, the person who made it, the moment it happened and where it came from — rules, libraries, mitigating controls, exceptions, violation dispositions, organisational scopes, connected systems and reference data alike. An auditor asking who loosened a rule, when and from what gets a direct answer.

WHERE KEYFORGE GOES FURTHER

The same rules, extended to the identities that aren't people

Service accounts and integration users already outnumber employees, and AI agents are arriving on top of them. Both hold real entitlements in your ERP. KeyForge governs them under **the same rules, the same severities and the same exception process** as everyone else — not a separate module with its own inventory and its own review cycle.

AGENT SEGREGATION OF DUTIES An agent's authority is evaluated, not assumed Every agent is registered with its vendor, its model and the tools it is entitled to call. Delegation is explicit — which human delegated authority, for what, and until when — and an agent whose delegator leaves the organisation loses its authority with them. Read-only agents are separated from those that can write or execute, because only the second group carries duty-segregation risk.	MCP SEGREGATION OF DUTIES A tool server is an entitlement bundle Grant an agent supplier creation on one MCP server and payment release on another, and you have built a toxic combination that exists in neither ERP. Each application sees one integration account behaving normally; the conflict lives in the composition of tools above them. KeyForge evaluates the agent's effective capability — every tool it can call, the identity it acts as, and everything it inherits by delegation — against the same rule library used for people.
RISK AT THE MCP LAYER Enforced at call time, not discovered afterwards The gateway checks the request as the agent makes it. A call that would complete a toxic combination is denied and recorded, with the rule that blocked it. That turns the agent layer from an unmonitored bypass around your ERP controls into an enforcement point that produces the same evidence as every other one.	JUST-IN-TIME CRITICAL ACCESS Elevated access for a purpose and a window Critical access is requested for a stated reason and a bounded period rather than held permanently — for people, for service accounts and for agents alike. Just-in-time eligibility and vault brokerage are tracked per account, sessions are logged and assessed, and use is reconciled against the reason given. Standing privilege becomes the exception that has to be justified.

Same rules. Same evidence. Same audit record — whether the requester is an employee, an integration or an agent.

The population that never leaves and is rarely owned

No joiner-mover-leaver event ever fires for a service account, ownership tends to walk out with the implementation partner, and the population is excluded from certification because nobody can attest to it. The register closes that gap.

Ownership	Exposure	Credential hygiene	Actual use
Named custodian, subject-matter expert and backup owner — so certification has someone to ask.	Environment, data classification, SOX scope and risk level, with the business service it serves.	Vault linkage, rotation interval, credential age and just-in-time eligibility, with overdue accounts surfaced.	Last-seen date and activity, so dormant accounts holding production entitlements are visible.

An integration account that can both create a supplier and release a payment is a toxic combination with no human attached — and it is analysed by the same engine, with the same severity, as if there were.

Runs alongside the identity platform you already own

KeyForge is the risk brain, not a replacement for your provisioning stack. It connects to applications read-only and returns decisions your existing tooling can act on.

Read-only from the source	Applications are read through catalogued extracts. Nothing is written back, so a risk assessment can start without a provisioning integration, a change window or a project to install anything in your ERP.
A decision your request workflow can call	The request-time conflict check answers in-line, so any access request system — your IGA suite, your service desk, or a home-grown workflow — can gate an approval on it without adopting a second rule matrix.
Findings pushed where your team already works	New high and critical findings are emitted as events, so they land in your monitoring, ticketing or alerting systems as they arise instead of waiting to be collected from another console.
Queryable by your own agents	The risk engine is exposed as a governed tool surface, so an approved assistant can pull a dashboard, trace a violation, run a what-if simulation or file an exception request — under the same permissions and audit trail as a person.

Three steps to a live picture of your estate

01	Scoping workshop Confirm the in-scope applications, the rule library to start from, and the integration and service accounts already in use.
02	Risk assessment Connect a read-only extract and return a live cross-application picture of duty-segregation and sensitive-access risk across the estate.

Walk the results, agree the remediation backlog, and size the programme against the roles and accounts actually generating the conflicts.