

AI IN THIRD-PARTY RISK MANAGEMENT

Meet the ProcessUnity TPRM AI Agents

Specialist agents built to manage key Third-Party Risk Management lifecycle processes — not to replace human judgment.

Table of Contents

Executive Summary	03
Part 1. Why a Single Generalist Agent Doesn't Work for Third-Party Risk Management	04
Part 2. What a "Built-for-TPRM" Agent Requires	05
A Hybrid Approach: Workflow Guardrails Plus Data Intelligence	05
Narrow Task Scope	05
Human-Gated Judgment, and Configurable by Design	06
Part 3. The Agents, One Cluster at a Time	07
Sourcing & Intake	08
Risk & Contracting	08
Due Diligence	09
Monitoring & Offboarding	10
Extending the Roster: Agent Architect and Bring Your Own AI	10
Part 4. How This Connects to the Rest of the ProcessUnity AI Stack	11
Conclusion: Bring ProcessUnity AI Agents to Your Program	12

WHO THIS IS FOR

Third-party risk, procurement, and cybersecurity leaders evaluating where AI belongs in a TPRM program.

WHAT YOU'LL GET

The full agent roster, what each one does, and where the analyst's judgment stays in the loop.

Executive Summary

Picture a Third-Party Risk Management (TPRM) program's strongest analyst on a Tuesday morning. They should be deciding whether a critical vendor's compensating controls are acceptable. Instead, they are on page 40 of a SOC 2 report, cross-referencing exceptions by hand. A new vendor request has been sitting in the intake queue for four days, and it turns out to be a duplicate of a supplier the business already uses, but nobody has had time to notice. A vendor's certificate of insurance expired last month and no one caught it. And the board deck is due Friday, so someone is about to spend three days pulling data out of half a dozen systems by hand.

None of that work requires the judgment a Third-Party Risk Management program exists to apply. All of it is manual labor that turns the program into a bottleneck that everyone in the business notices. New vendor onboarding takes weeks, findings go quiet, and the people trusted for their judgment spend days on work that never needed judgment at all.

The market's answer is usually a single, broad AI agent asked to handle third-party risk. **The truly impactful AI solutions offer something different: a roster of task-level specialists.**

Each one is built for one job in the lifecycle, each grounded in real risk data, and each leaving the judgment-bearing call to the analyst. ProcessUnity TPRM AI Agents follow this line of operation. Let's break down what each AI agent does, what decisions stay with the analyst, and how the whole approach fits together.

NARROW BY DESIGN

Every agent does one job an analyst can state in a sentence.

GROUNDING IN RISK DATA

Program history and Exchange-attested evidence, not open-web guesswork.

HUMAN-GATED

Consequential decisions stay with the analyst, and stay traceable.

01

PART ONE

Why a Single Generalist Agent Doesn't Work for Third-Party Risk Management

The market's default answer to "agentic AI" in Third-Party Risk Management is a chatbot bolted onto an existing platform, or one broad agent stretched across the entire lifecycle. Both versions share the same underlying weakness: their scope is too large to state in a sentence, which means it is also too large to verify, govern, or defend when an auditor asks what, exactly, the AI was allowed to decide.

Third-party risk work resists this kind of generalist treatment because it is not one task, but instead is a portfolio of dozens of distinct analyst jobs, each with its own regulatory nuance and its own tolerance for error. What counts as sufficient evidence for an encryption control is a program-specific standard. What DORA's ICT third-party provisions require of a register is not what APRA CPS 230 expects of a service-provider inventory. A single model asked to "do" third-party risk has to hold all that context at once, for every task, which is exactly the condition under which a generalist tool starts producing output that sounds right, but is not.

ONE GENERALIST AGENT

Scope too broad to state, verify, or defend.
Holds every program's context at once, for every task.

A ROSTER OF SPECIALISTS

Each agent is narrow enough to describe in a single sentence, and meant for consistent verification.

The alternative is not less AI. It is AI shaped to the size of the job: a roster of specialists, each one narrow enough that an analyst can say, in a sentence, what it does, and can verify, every time, that it did it.

SEE IT WITH YOUR OWN DATA

Request a Demo of TPRM AI Agents

Watch a ProcessUnity practitioner run the agents live — a duplicate check at intake, a SOC 2 broken down to its exceptions — and see exactly where the automated and gated boundaries would sit in your program.

[REQUEST A DEMO](#)

What a "Built-for-TPRM" Agent Requires

Three requirements separate an agent that a risk program can rely on from one it will eventually have to walk back. They form the recurring thread behind every ProcessUnity TPRM AI Agent introduced throughout this paper.

01 A Hybrid Approach: Workflow Guardrails Plus Data Intelligence

An agent built for Third-Party Risk Management runs within a hybrid model, governed by deterministic workflow guardrails and powered by TPRM data intelligence, rather than left to reason from scratch on every run. For ProcessUnity TPRM AI Agents, the guardrails come from the TPRM Platform — the same rules, routing logic, and escalation paths that already orchestrate a program's processes. The intelligence comes from data most general-purpose tools never see: a program's own historical records, plus the attested assessments and evidence available through the ProcessUnity Global Risk Exchange.

REQUIRES A platform, like ProcessUnity, that already manages a program's workflow logic, and an agent design process shaped by practitioners who know what a given piece of evidence is supposed to prove, not just a capable model.

THE OUTCOME An agent that acts like it is part of the team, not a text generator with no memory of the program it is working inside.

02 Narrow Task Scope

Each agent in the ProcessUnity roster is built to do one job, with defined inputs and outputs. A "duplicate-vendor check" is a narrow, bounded task. "Handle vendor risk" is not. The narrower the task, the easier it is for an analyst to state exactly what the agent touches, and the easier it is to verify the agent got it right. Accuracy verification is what will earn an AI capability trust over time.

REQUIRES A task definition specific enough to fit in one sentence, with inputs and outputs an analyst can name without hedging.

THE OUTCOME An agent whose full scope of action is never a mystery, to the analyst using it or the auditor reviewing it.

03 Human-Gated Judgment, and Configurable by Design

The third requirement is that agents are fully automated where the task is deterministic — a duplicate check, or an entity lookup — and gated behind human approval wherever judgment is required for the outcome. That boundary is not a limitation bolted on after the fact, but configurable by design with ProcessUnity. The degree of human-in-the-loop engagement on any given agent can be set based on how the program wants that agent to handle ambiguity, present results, and take next steps, so the boundary between automated and gated reflects the program's own risk appetite rather than a fixed default.

REQUIRES A clear designation, for every agent, of what runs automated and what runs gated, and the ability to adjust that designation as a program's governance matures.

THE OUTCOME An analyst who approves every consequential risk decision, and a program that can show exactly which decisions those were, even if they were part of an agent workflow.

WHERE THE BOUNDARY SITS

Every agent in the roster lands in one of three places, and a program decides which.

RUNS AUTOMATED

Deterministic checks with a verifiable answer: duplicate records, entity lookups, coverage validation, contract term extraction.

**REMOVE
ROADBLOCKS**

RUNS GATED

Anything carrying a judgment: control sufficiency, risk acceptance, remediation adequacy, executive narrative.

**ANALYST
APPROVES**

CONFIGURABLE

Programs move the line themselves as governance matures — the default is a starting point, not a ceiling.

**CONTINUED
GROWTH**

The Agents, One Cluster at a Time

The [ProcessUnity TPRM AI roster](#) is organized around where work breaks down in the vendor risk management lifecycle, including sourcing and intake, risk and contracting, due diligence, and monitoring and offboarding. Each agent hands back a structured, attributed result to the requested task, the same way a good colleague hands over a finished piece of work, not a rough draft.

Below are the ProcessUnity TPRM AI Agents, each with their own determined task to reduce manual work and promote efficiency across the Third-Party Risk Management lifecycle at each step.

01

SOURCING & INTAKE

2 AGENTS

Duplicate Vendor Inspector

Duplicate Service Inspector

02

RISK & CONTRACTING

2 AGENTS

IRQ Responder

Insurance Guardian

03

DUE DILIGENCE

4 AGENTS

SOC 2 Analyzer

Trust Center Finder

Contract Scanner

Entity Identifier

04

MONITORING & OFFBOARDING

2 AGENTS

Executive Risk Reporter

Issue Remediator

The current list of TPRM AI Agents as of fall 2026. ProcessUnity's roster of agents will develop in future months.

■ Sourcing & Intake

The moment a new vendor or service request enters the pipeline is also the moment redundant work is the cheapest to catch, because it's before a record exists, before a contract has been drafted, and before anyone has spent time on it.



Duplicate Vendor Inspector

Flags duplicate third-party records at intake, before a new vendor record is even created.



Duplicate Service Inspector

Surfaces redundant spend when a requested service is already provided or under contract elsewhere in the business.

REQUIRES A view across the program's existing vendor and service records instantly when a new request comes in.

THE OUTCOME Vendor requests that duplicate existing coverage get caught in seconds, not discovered months into onboarding.

■ Risk & Contracting

Once a request is confirmed as genuinely new, the next TPRM bottleneck is the paperwork that must happen before a vendor relationship can move forward: an inherent risk questionnaire needs to be completed, or insurance coverage must be confirmed. ProcessUnity has AI Agents to support these steps.



IRQ Responder

Pre-fills the inherent risk questionnaire from available vendor intelligence so the team only needs to confirm risk exceptions instead of entering data into every field.



Insurance Guardian

Validates certificates of insurance against required coverages, limits, and dates at the moment they are submitted, not months later at renewal.

REQUIRES Structured vendor intelligence and a program's own coverage thresholds to check against.

THE OUTCOME Intake paperwork that used to take days of back-and-forth clears in the time it takes to review one exception.

■ Due Diligence

Due diligence is where the heaviest manual burden in the lifecycle has always lived, because it includes reading pages and pages of vendor evidence, chasing documentation, and confirming who a vendor actually is before anyone assesses them. Here's how ProcessUnity AI Agents relieve that burden.



SOC 2 Analyzer

Extracts scope, exceptions, and complementary user entity controls (CUECs) from a SOC 2 report, mapped to the program's own controls.



Trust Center Finder

Finds and harvests a vendor's published trust center evidence, ending the manual document chase.



Contract Scanner

Pulls the terms the program depends on out of an executed agreement — dates, owner, value, jurisdiction, renewal mechanics, and service levels — each written back with its source clause.



Entity Identifier

Confirms a vendor's legal entity against the GLEIF registry — Legal Entity Identifier, registered name, jurisdiction, and status — before due diligence begins.

REQUIRES Access to the vendor's submitted evidence, published trust center data, executed contracts, and the GLEIF entity registry.

THE OUTCOME An analyst only needs to read a confirmed exception list instead of a 90-page PDF, and starts due diligence already knowing who they're dealing with.

■ Monitoring & Offboarding

The lifecycle does not end with onboarding and due diligence assessments. Portfolios need to be explained to executives on demand, and findings need to continuously move forward instead of stalling in a queue.



Executive Risk Reporter

Drafts executive risk briefs and one-click closeout summaries on demand, turning sprawling portfolio data into something actionable with a five-minute read.



Issue Remediator

Drafts remediation plans for open findings and routes the resulting actions to their owners.

REQUIRES Live portfolio and findings data, plus a program's own remediation language and escalation paths.

THE OUTCOME A one-page brief ready before the meeting starts, and findings that move within days instead of dragging for weeks.

Extending the Roster: Agent Architect and Bring Your Own AI

The four clusters above cover where the current ProcessUnity AI Agents work, but the roster is not fixed.



Agent Architect

The no-code way to build new agents and modify any agent already in place, letting teams configure prompts, tools, and guardrails without a development cycle. When a program identifies its own high-value, narrow task that isn't covered yet, building the agent becomes a configuration exercise rather than a procurement decision.

NO DEVELOPMENT CYCLE



Bring Your Own AI

BYOAI from ProcessUnity extends the same flexibility to the model layer. Any agent can run on a program's own model licensing agreement, bringing existing model investment into the platform rather than asking a program to standardize on a single vendor's model choice.

YOUR MODEL, YOUR TERMS

Both of these flexible additions allow TPRM teams to incorporate the value of ProcessUnity AI Agents in a way that works for their unique program needs, on their own time.

04

PART FOUR

How This Connects to the Rest of the ProcessUnity AI Stack

The agents introduced above sit alongside AI capabilities that have been in production longer and take a different approach to grounding: Evidence Evaluator and Assessment Autofill are trained directly on ProcessUnity's own TPRM dataset, and are built to support risk and client assurance teams with time-consuming assessment steps.

Evidence Evaluator

Reads SOC 2 reports, ISO certifications, and security policies the way an experienced internal risk analyst would, citing exactly where in the document a control is or is not supported — turning a multi-hour document review into a confirm-or-flag step.

Assessment Autofill

Pre-populates a third party's questionnaire responses from their own prior submissions, cutting average vendor response time by over 90%.

90% faster vendor response time

Together, the agents and these established tools form one continuous approach rather than two separate product lines, shaped by the same decade of accumulated practitioner expertise and backed by the same governance standards — every run is logged and traceable, every judgment-bearing decision gated behind a human, and no output the program can't explain to the auditor who asks about it.

IN PRODUCTION AT A LARGE GLOBAL TECH FIRM

75%

fewer returned IRQs

54%

faster intake

43%

higher throughput

45%

of assessments fully agentic

CONCLUSION

Bring ProcessUnity AI Agents to Your Program

A single generalist agent asked to "handle third-party risk" will always be either too shallow to trust on the hard calls or too broad to verify on the easy ones.

The roster approach solves both problems at once: specialists narrow enough to bound and verify, grounded in real risk data, configurable to a program's own governance appetite, and expandable through Agent Architect the moment a new high-value task shows up.

Hours back per assessment, days instead of weeks to onboard a new vendor, and **an analyst spending their time on the calls that actually needed their judgment.**

To get started working with ProcessUnity AI Agents in your program today, contact the ProcessUnity team or visit our website for more information and resources.

See the agents in action

Walk through the roster with a ProcessUnity practitioner and see where the automated and gated boundaries would sit in your program.

REQUEST A DEMO

www.processunity.com



ABOUT PROCESSUNITY

ProcessUnity is *the* Third-Party Risk Management (TPRM) company. Our software platforms and data services protect customers from cybersecurity threats, breaches, and outages that originate from their ever-growing ecosystem of business partners. By combining the world's largest third-party risk data exchange, the leading TPRM workflow platform, and powerful artificial intelligence, ProcessUnity extends third-party risk, procurement, and cybersecurity teams so they can cover their entire vendor portfolio. With ProcessUnity, organizations of all sizes reduce assessment work while improving quality, securing intellectual property and customer data so business operations continue to operate uninterrupted.

See how at www.processunity.com.

ADDRESS

ProcessUnity
33 Bradford Street
Concord, MA 01742
United States

SOCIALS

Twitter: [@processunity](https://twitter.com/processunity)
LinkedIn: [processunity](https://www.linkedin.com/company/processunity)

CONTACT

www.processunity.com
info@processunity.com