ILLUSTRATIVE CASE STUDY / SOUTH AFRICA

Preventing AI data leakage

An illustrative case study of Acronis GenAI Protection

AI becomes more useful when employees supply context. That same context can include the information a business is responsible for protecting. Controls need to act where people share data, with an owner responsible for the decisions and the evidence.

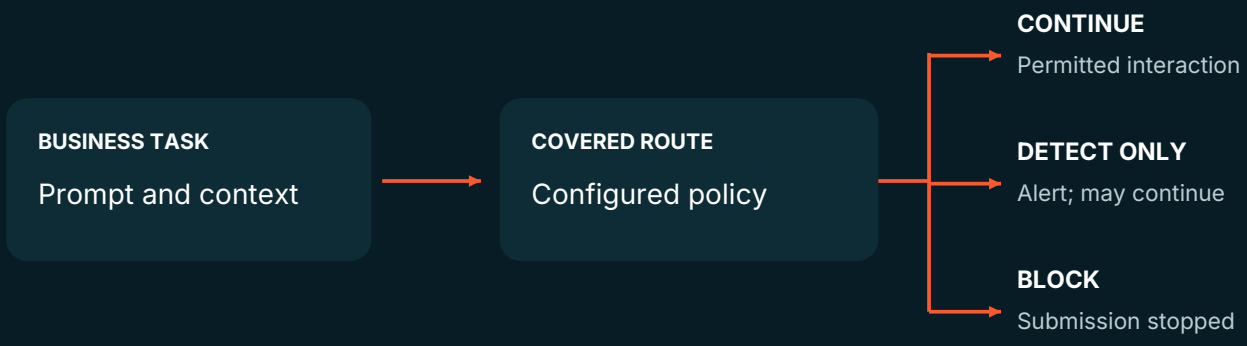

A practical control model, a proposed pilot and the evidence a business should ask for.

Illustrative business scenario. The company, workforce and events are assumptions. No customer deployment, testimonial, measured saving or breach reduction is claimed.

01 THE SCENARIO

A better proposal can create an unwanted disclosure

A 60-person South African advisory business wants staff to prepare proposals and respond to customers faster. For this scenario, employees use managed Windows laptops; some also work on personal devices. AI use has grown through individual choices, without a consistent register of approved tools or permitted data.

A sales employee plans to paste customer notes into an AI assistant to improve a proposal. The notes include contact details, negotiated pricing and internal comments. The task has a legitimate business purpose. The chosen account and data transfer still need approval.

Workflow	Information entering the decision	Intended business value
Sales	Proposal notes and customer contacts	A polished proposal with less drafting effort
Support	Diagnostic extracts that may contain secrets	A clearer explanation of a technical issue
People management	Employee correspondence	A concise summary of a sensitive matter
Marketing	Public information and approved messaging	Faster preparation of campaign copy

These workflows deserve different rules. Public marketing material may be suitable for an approved service. Personal information, secrets and confidential commercial details require a more restrictive decision. OWASP identifies sensitive information disclosure as a risk in LLM applications. [2]

The leadership question is whether staff can complete valuable work without making unauthorised data-sharing decisions. Potential consequences include investigation effort, customer concern and interruption of the workflow; this scenario does not assign an invented financial loss.

The supporting desk-research workbook catalogues 300 AI offerings across 28 categories. It illustrates market breadth, not AI adoption rates, comparative safety or an Acronis compatibility list.

02 THE CONTROL MODEL

Protect the decision before information is submitted

The proposed design combines an approved-use standard with controls on covered AI interactions. The workflow owner defines what is allowed; the technical team validates how the configured policy behaves.

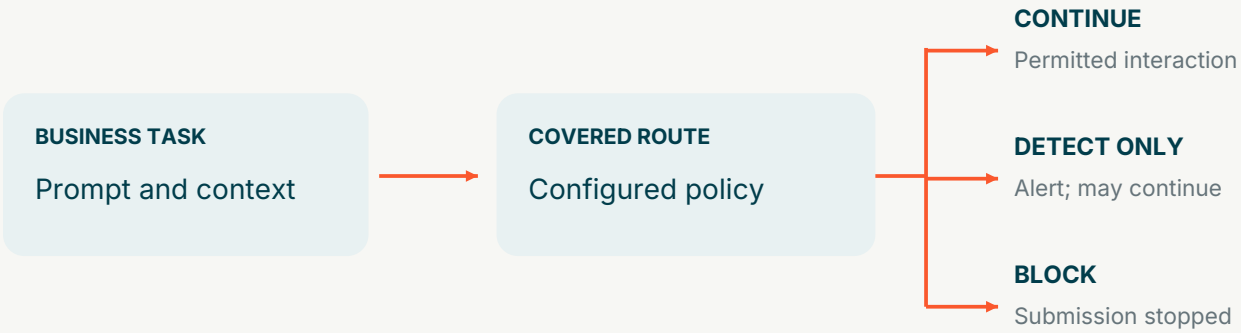

Conceptual policy outcomes on a covered route. Detection and enforcement depend on the configured mode and the rules that apply.

Make usage visible

Use reporting to establish where review is needed. Assign an owner to each approved workflow and review the evidence regularly. [4]

Set the application policy

July 2026 documentation adds application and domain allow/block policies. Confirm the policy deployed in the target environment. [5]

Test the prompt rules

The data sheet documents sensitive-data classification, harmful-prompt controls and detect-only or block modes. Use synthetic material to validate the selected rules. [6]

Act on the event

Acronis documents reviewing blocked transfer attempts in the events log. Define who investigates, coaches the user and approves an exception. [3]

Return to the proposal

For the proposal task, test both a prohibited version containing policy-recognised sensitive material and a permitted, sanitised version. The intended result is to stop the prohibited submission on the covered route while preserving the useful drafting task. This is a test objective, not an observed result.

03 SCOPE AND EVIDENCE

Access control and content inspection are different tests

A credible deployment records what each control actually covers. Acronis positions prompt protection around browser-based AI interactions; its July release also describes application access control across web, desktop and SaaS scenarios. That broader access scope does not establish inspection of every payload. [1,5]

Risk or decision	Relevant control	Evidence to collect
Unapproved application	Application or domain allow/block policy	Does the intended policy block the selected access route?
Sensitive prompt	Prompt DLP on a covered interaction	Is prohibited synthetic content blocked, with legitimate work still usable?
Harmful prompt	Documented harmful-prompt detection	Does the agreed test trigger the expected action and event?
Investigation	Events, widgets and user context	Can a reviewer link the event, logged-in account, device and policy?
Files, voice, APIs and personal devices	Coverage must be established separately	What control applies to each path, and what remains uncovered?

For broader local and network data channels, Acronis points to its wider DLP service. [1] Do not assume that a browser policy controls a custom API agent, native application payload or unmanaged device. Validate those routes and record the additional controls required.

Prompt injection can also arrive through external content. OWASP recommends restricted privileges and approval for consequential actions. [8] A prompt security feature does not replace secure agent design, provider assurance or checks on AI-generated answers.

Confirm current OS and browser support, agent version, supported AI destinations, licensing and tenant availability. The February launch notes provide a dated Windows support baseline; they are not a substitute for the current deployment check. [7]

04 THE PROPOSED PILOT

Build evidence before claiming an outcome

A four-week pilot gives the business a manageable starting point. This sequence is a Soteria editorial recommendation, not a vendor implementation commitment. Use synthetic test data and keep real sensitive work on an approved route.

WEEK 1 **Define** Business owners record permitted tasks and data. IT confirms coverage and prerequisites. Explain monitoring, access to records and retention to staff.

WEEK 2 **Observe** Establish a limited baseline with detect-only mode where appropriate. Observation is not enforcement; maintain other controls on sensitive work. [6]

WEEK 3 **Validate** Test permitted and prohibited prompts, access policies, user attribution and exceptions. Include representative local data formats; retain policy versions and results.

WEEK 4 **Enforce and review** Apply agreed blocking to validated workflows. Retest changes, resolve false positives and assign owners to the remaining coverage gaps.

Measure the control and the work

Measure	What to record
Coverage	Tested in-scope routes / defined pilot routes
Control behaviour	Correctly blocked prohibited tests and permitted legitimate tests
Operational effort	Investigation time and exception resolution time
Usability	False positives and user completion of approved tasks
Business value	Drafting time alongside human review and correction effort

An event count alone is not a success measure. More events may mean better visibility; fewer may mean less use or missing telemetry. Report the tested scope, the limitations and the effect on work together.

The business owner approves purpose and data. IT or the MSP operates the controls. The Information Officer or privacy lead reviews personal-information handling. Leadership owns the residual risk and the decision to expand.

05 THE LEADERSHIP LESSON

Make safe AI use an operating responsibility

For organisations that allow external AI, enforceable control over sensitive data belongs in the operating baseline. Acronis GenAI Protection can support that baseline within an established platform; the proposed pilot makes its suitability and limitations visible.

The intended change is practical: employees know the approved route, prohibited sharing can trigger a configured control on covered interactions, and someone owns the resulting review. This case reports no achieved customer outcome. It defines what a real deployment would need to demonstrate.

POPIA section 19 addresses reasonable technical and organisational safeguards; section 72 addresses transfers outside South Africa. [9] Controls can support these responsibilities alongside appropriate terms, purpose, minimisation and review. The product does not guarantee compliance. This is general information, not legal advice.

Discuss an AI usage and protection review

Start with one workflow that matters. Ask Soteria Cloud to help scope an AI usage and protection review, define the coverage questions and agree a practical pilot.

web.soteriacloud.com/contact

Soteria Cloud is a trusted Acronis partner and cloud aggregator supporting South African MSPs and businesses.

Sources and research basis

Research checked 17 September 2026. Product statements are based on published Acronis material, not independent effectiveness testing. The scenario, pilot and business interpretation are original editorial analysis.

1. Acronis product overview and scope
2. OWASP Sensitive Information Disclosure
3. Acronis GenAI DLP demonstration
4. Acronis monitoring and reporting demonstration
5. Cyber Protect Cloud 26.07 release notes
6. Acronis GenAI Protection data sheet, May 2026
7. Cyber Protect Cloud 26.02 release notes
8. OWASP Prompt Injection
9. POPIA, sections 19 and 72