bonfy.ai

AI-Native Data Security:
What Leaders Must Rethink as Copilots and Agents Reshape Work

Jacob Bardin

Enterprise AI has become a consumer of enterprise knowledge

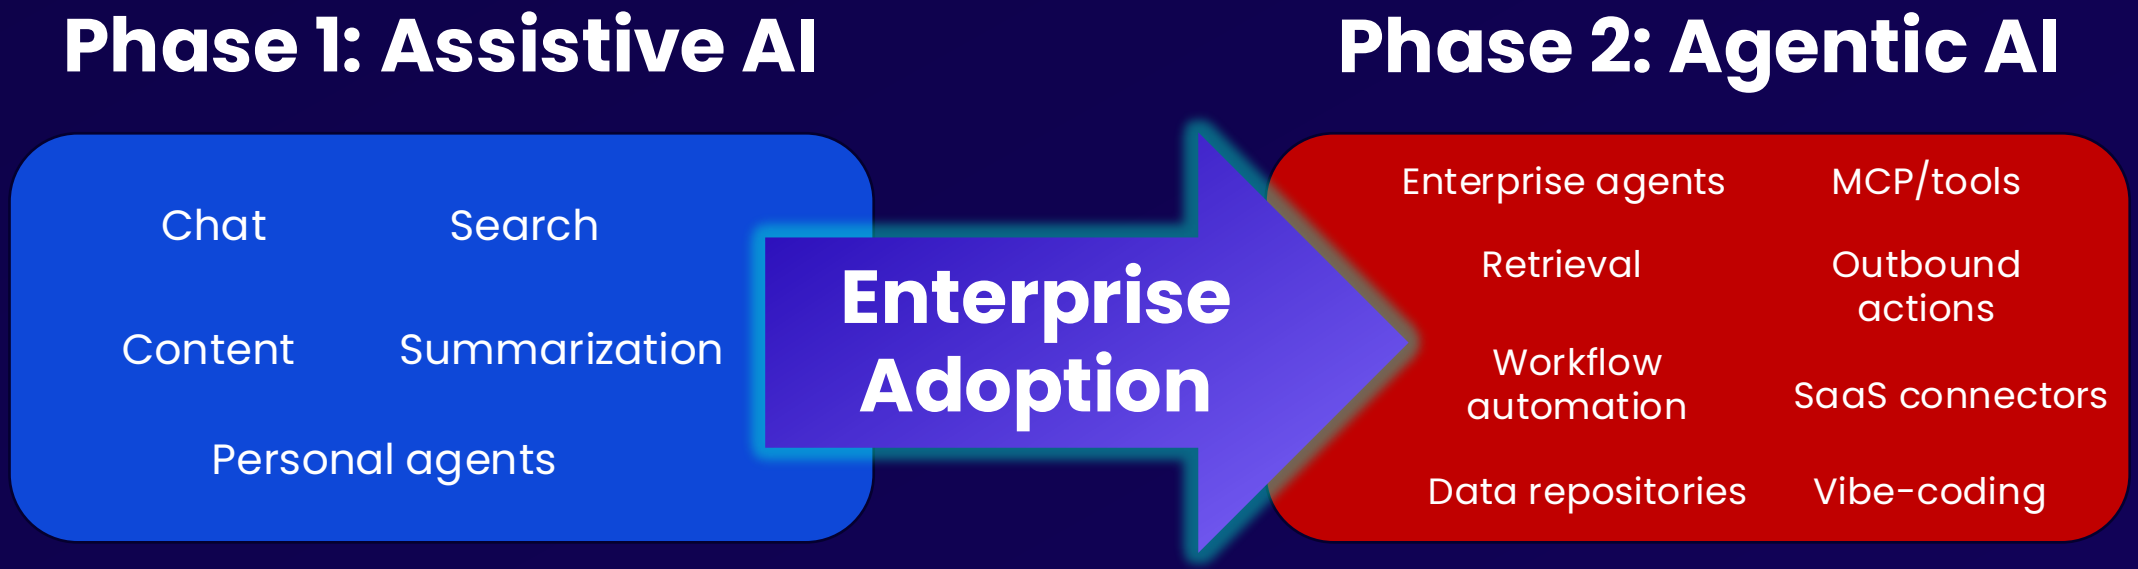

AI becomes valuable when it connects to enterprise data

That is exactly when the **security model breaks**

AI risk is expanding, while humans are getting out of the loop

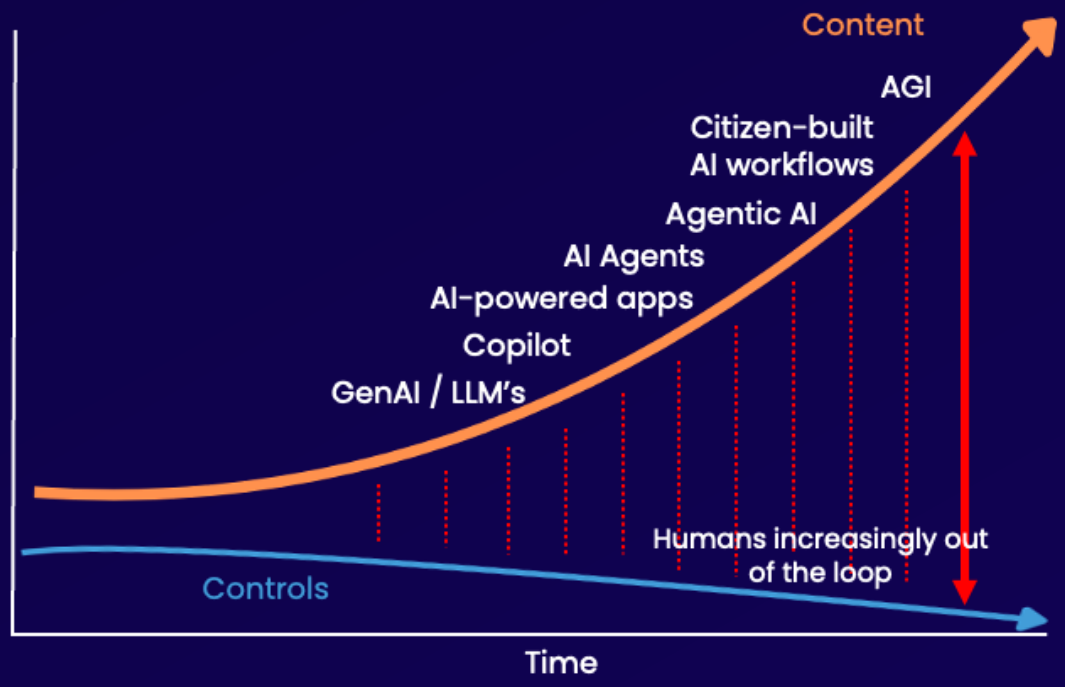

A copilot summarizes an email thread and surfaces details that were never meant for broad visibility.

An employee uses an external AI tool to analyze content that includes sensitive business data.

An AI agent combines and sends customer-specific information to the wrong customer.

Without human context, sensitive data is being used in ways organizations can't control

What actually happens after an agent gets the command

User prompt:

"Summarize open renewal risks for Acme and email me the key issues"

Retrieves enterprise data the user can access

Pulls, combines, and surfaces sensitive information

Interprets, summarizes, transforms, or stores in memory

Calls tools, draft messages, update records, or share content

Response

The response looks clean and helpful.

Goes out in email and chat.

Ignored security layer: Live data movement during processing

Risk callouts

- Over-retrieval
- Sensitive combination
- Memory persistence
- Tool misuse
- Wrong-channel output

What's needed:
Context-aware control layer that fills the gaps humans leave behind

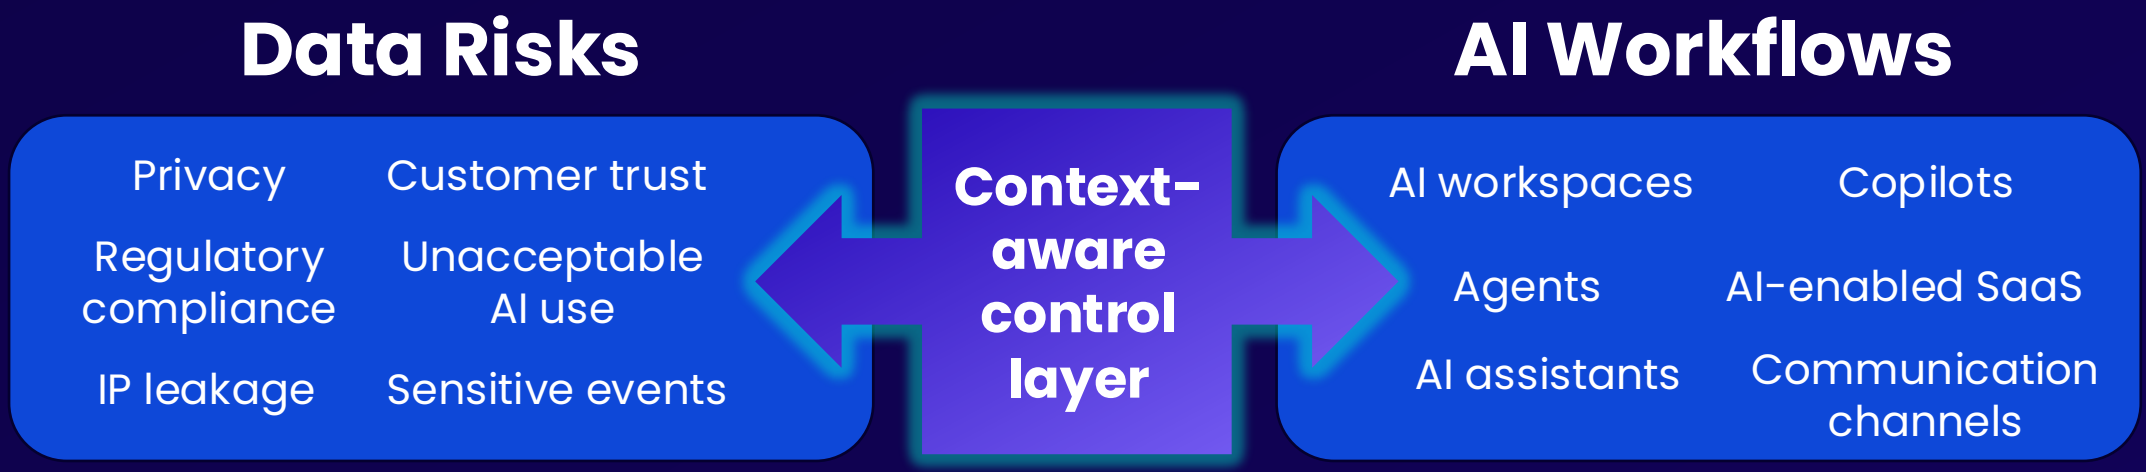

"Build a proprietary context-driven platform for trust, risk and security by deploying an AI context layer that enables AI automation with trust."

-**Gartner**; AI Context Layer Wins the New AI Security Race; March 31 2026

Enterprise AI requires runtime, contextual policy decisions

Every decision must be multi-dimensional

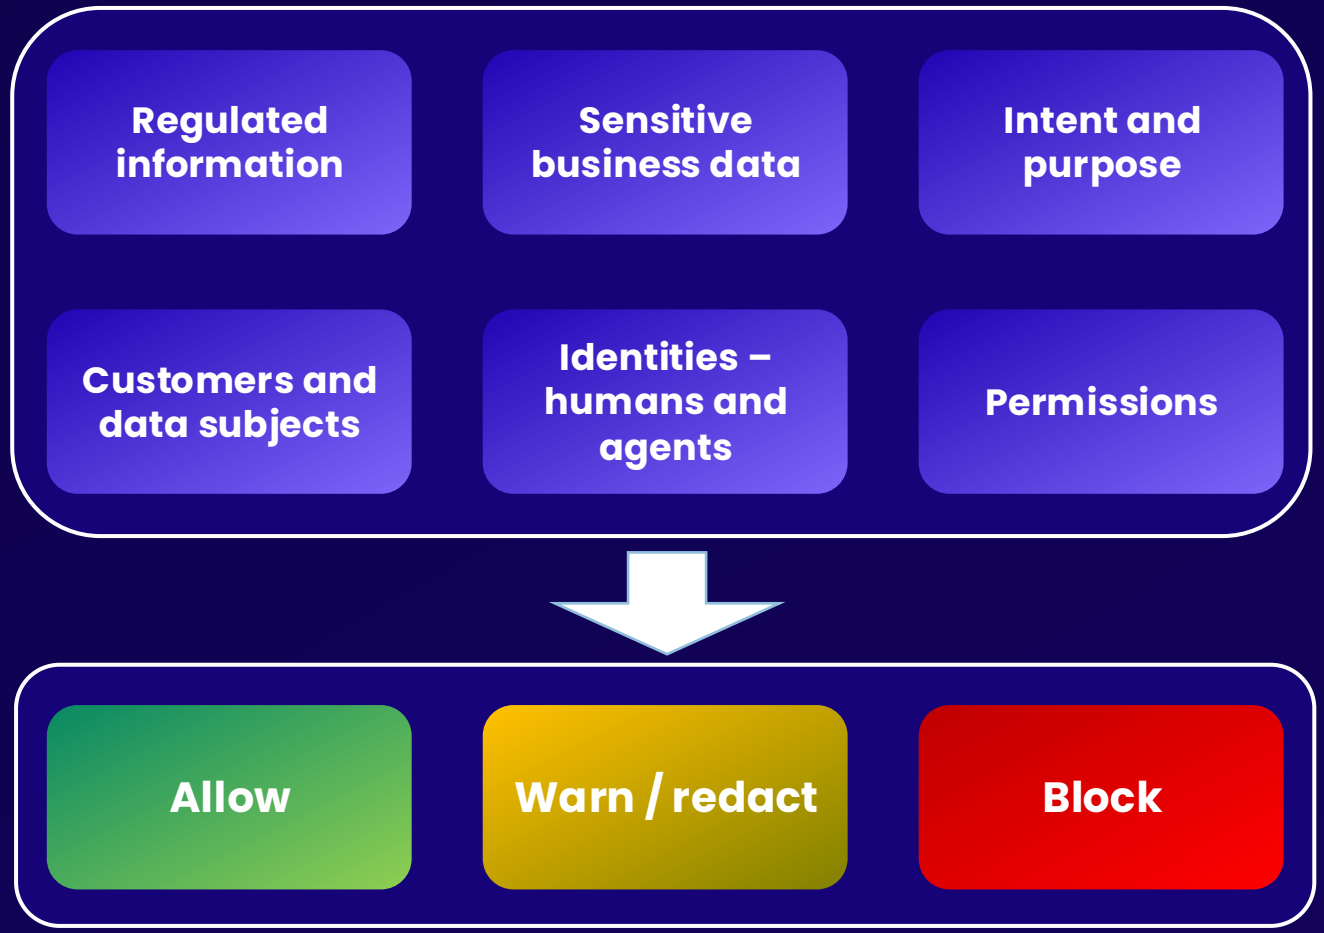

What control layer for AI data must do

01

Understand the data

Classify sensitive content, detect forbidden topics, and identify policy-relevant signals.

02

Understand the business relationship & context

Link people, customers, agents, topics, and permissions to the content being used.

03

Enforce in real time

Apply guardrails before, during, and after AI use across prompts, retrieval, generation, and sharing.

The goal is not to block AI; it is to make AI use governable, auditable, and safe at enterprise scale.

Without accuracy, enforcement fails

Accuracy = detecting real risk without noise or missed violations

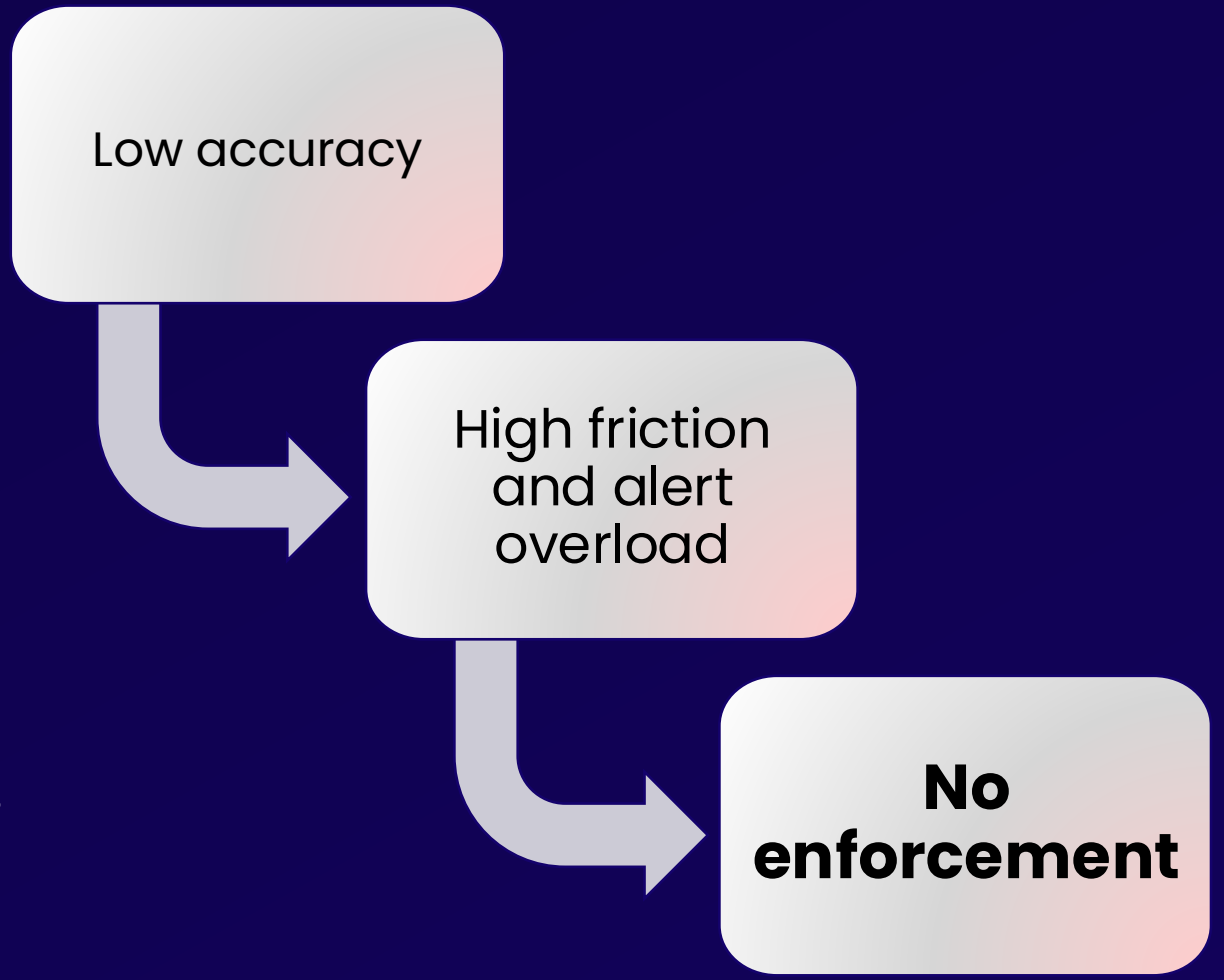

Holistic control layer for 2D data flows

Data visibility & control

West–East

Across systems and data flows

North–South

Across users, agents, reasoning, tool execution

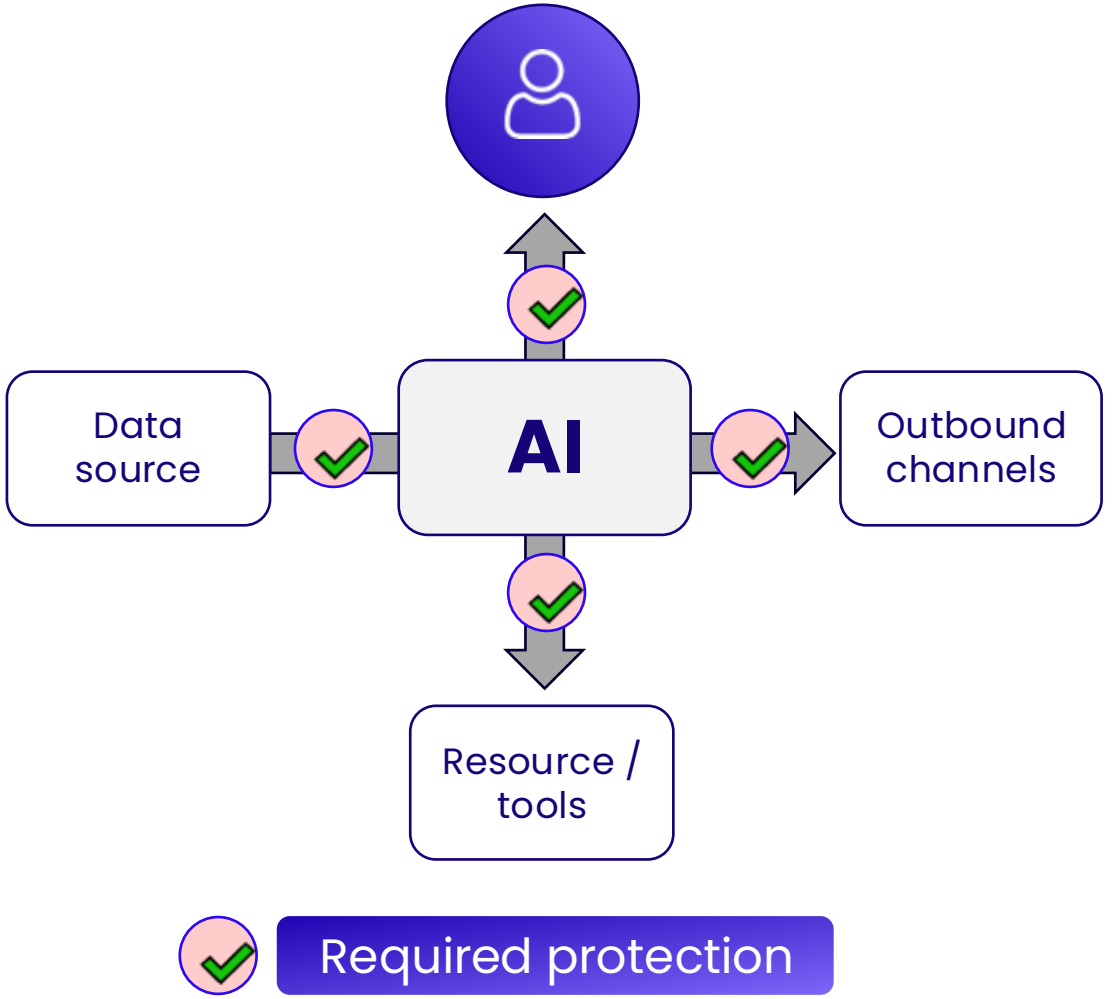

Three Control Planes of Enterprise AI

Configuration

What the agent is allowed to do, what tools it can call, and how the workflow is set up.

Identity

Who or what is acting, what privileges it has, and which trust boundaries it crosses.

Data

What sensitive content is entering the workflow, how it is being used, what context it carries, and where it is going next.

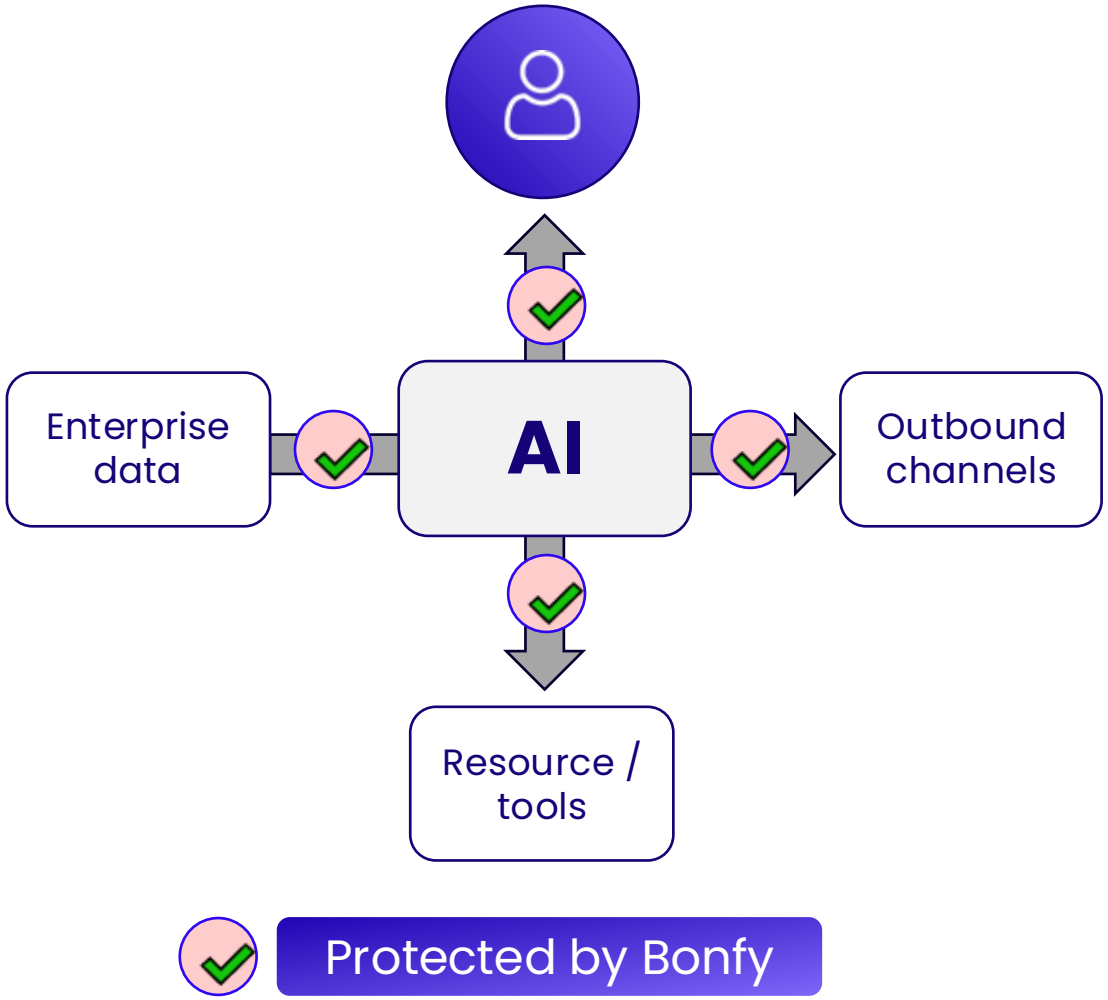

Operationalizing AI Governance

Data

What data can be accessed, transformed, or shared?

Application

What AI systems, copilots, and agents are allowed to use it?

Identity

What are what is acting, and what authority should they have?

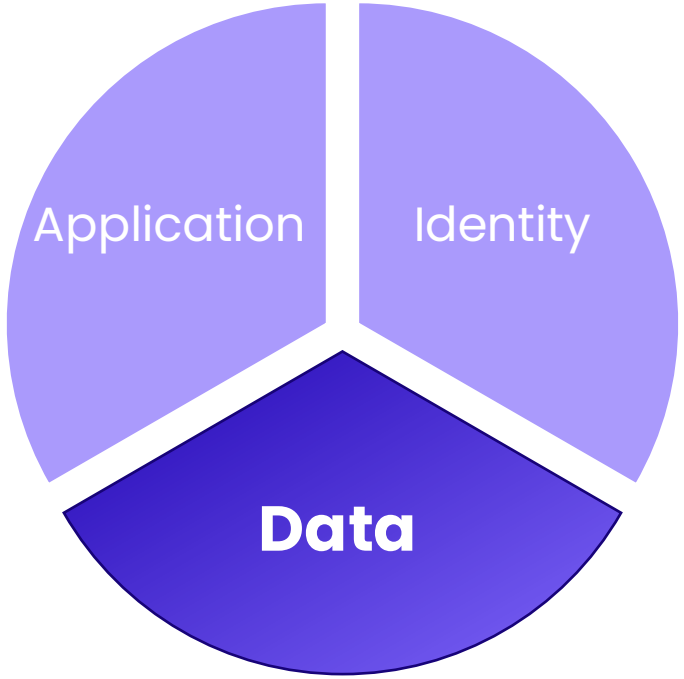

If you cannot control how sensitive data is used in real time, you do not yet have effective AI governance.

- Apply policy where data is used, not only where it is stored
- Govern the full lifecycle: input, in-use, and output
- Make human and non-human identity part of every decision

How to operationalize AI governance

If you can't control how data is used, you don't have AI governance

- **Data** — what data can be used, and how
- **Application** — what AI systems and agents can do
- **Identity** — who (human or agent) can act

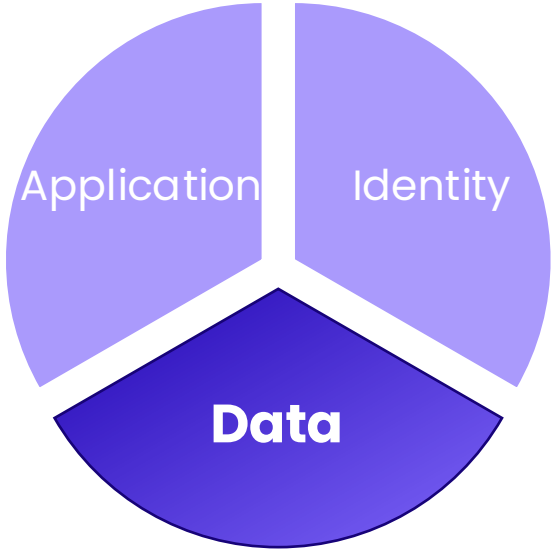

AI governance ultimately depends on the data layer

- Enforce the intended use of data
- Control at the moment data is used
- Across the lifecycle: **Input** → **In-use** → **Output**

Three questions every executive team should ask now

1. Where is AI already interacting with sensitive data across the business?

2. Which copilots, agents, and embedded AI features operate outside approved controls?

3. Can policy be enforced at the moment data is used, not just after exposure?

Thank you

Bonfy.ai

/bonfy-ai

jacob.bardin@bonfy.ai

bonfy.ai

Business context is critical for high accuracy

Example: Raul is using **ChatGPT** / **Claude Chat** to write a message intended for Alex Taylor (CFO) and John Smirnov (VP)

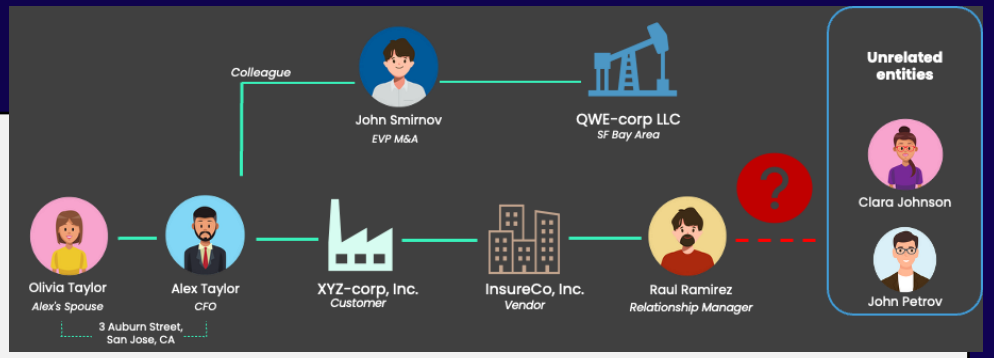

To: alex.taylor@xyz-corp.com

CC: john.pertrov@outlook.com

Subject: Ref 911211164

What traditional DLP finds?

or your account information, here is th

Data leakage

Saving account 44721211098: balance $412,900,114.00

Regulatory exposure

I would like to confirm that the policyholder list was sent to John Smirnov (in CC) for the cancellation notice effective next week.

BTW, John, do you still live on 3 Auburn Street, San Jose, CA? We need it for the forms.

Best wishes to your wife, Clara,

Raul Ramire

Enterprise S

Record integrity violation

Privacy violation

Business context is about **entities** and their **relationships** - employees, customers, partners, agents