



A GLOBAL IR GROUP FIELD REPORT · 2026

# The State of *AI* in Investor Relations

What IROs are actually using, what they refuse to automate, why most pilots quietly fail, and what an AI-native IR operating model looks like when it works.

---

Written IRO-to-IRO.

Global IR Group · Fractional IR Leadership & AI-Enabled Capital Markets Intelligence

2026 EDITION  
VOL. 01

## THE PREMISE

# The question changed when no one was looking.

Two years ago, the question in every IR department was “*should we be using AI?*” In 2026, that question is settled. Everyone is using it. The honest question now is quieter and far more uncomfortable: “*why isn’t ours working the way the demo promised?*”

Investor relations is one of the most exposed functions in the company to get this right. The output is read by analysts, quoted in filings, and governed by Regulation FD. The team is small. The calendar is unforgiving. And the cost of a confident, well-formatted mistake is measured not in embarrassment but in disclosure risk.

This report is a practitioner’s read on where AI in IR actually stands: drawn from the field, not from a vendor deck. It covers what IROs are really using, the fears that keep adoption shallow, why most pilots stall, and the discipline that separates the programs pulling ahead from the ones generating impressive-looking noise. It closes with the model we believe the next decade of IR will be built on.

## THE THROUGH-LINE

AI has not changed what investor relations is. It has changed how much of the *work* an IRO can hand off, and made the IRO’s judgment more valuable, not less.

## INSIDE THIS REPORT

- |    |   |  |
|----|---|--|
| 01 | <b>What IROs Are Actually Using</b>           | The reality behind the adoption headlines          |
| 02 | <b>The Three Fears</b>                        | Accuracy, governance, and the second full-time job |
| 03 | <b>Why Most AI Pilots Fail</b>                | The predictable anatomy of a stalled rollout       |
| 04 | <b>What Top IROs Automate, and Never Will</b> | The line the best practitioners refuse to cross    |
| 05 | <b>The Future of IR Operating Systems</b>     | From tools you open to a layer that runs           |
| 06 | <b>The AI-Native IR Operating Model</b>       | What it looks like when it actually works          |

SECTION 01

# What IROs are actually using

Strip away the conference panels and the platform announcements, and the real picture is more modest, and more revealing, than the headlines suggest.

The dominant tool in investor relations today is not a purpose-built IR platform. It is a general-purpose chatbot open in a browser tab. IROs reach for ChatGPT, Claude, and Microsoft Copilot the way they once reached for a search engine: constantly, informally, and almost always for the same handful of jobs.

## Drafting

First passes of earnings scripts, press releases, Q&A trees, and bios. The blank page is gone; the editing is not.

## Synthesis

Summarizing analyst notes, condensing earnings transcripts, and digesting peer filings into something readable before a meeting.

## Sense-making

“Explain this competitor’s quarter,” “reframe this for a generalist PM,” “what would an analyst push on here?”

This is real value, and it is widespread. But notice what it is: **personal productivity, used ad hoc, by individuals, with no shared system underneath.** The IRO who summarizes a transcript today starts from scratch tomorrow. Nothing accumulates. The intelligence lives in a chat window and evaporates when the tab closes.

## The shadow-AI pattern

Because adoption has run ahead of policy, most IR teams are living with what security teams call “shadow AI”: powerful tools used productively but unofficially, often pasted with material drafts and non-public data, with no record of what went where. The function is more AI-enabled than its leadership realizes, and less governed than its risk profile demands.

## Most

AI use in IR today is general-purpose chat, not a dedicated IR system

## Few

Teams have a written policy governing what may be entered into a model

## Fewer

Have a system where prior work, context, and sources actually persist

### THE TAKEAWAY

The gap in IR is no longer access to AI. It is the gap between *a tab you visit* and *a system that remembers*, carrying your filings, your investors, your peers, and your history from one task to the next.

## SECTION 02

# The biggest fears

Ask a room of IROs what holds them back, and the answers collapse into three. None of them are irrational. All of them are addressable, but only by design, not by enthusiasm.

## Fear 1: Accuracy

In most functions, an AI hallucination is an inconvenience. In investor relations, a confidently wrong number in an earnings script, a guidance figure, or an analyst response is a potential **disclosure event**. The fear is not that the model is occasionally wrong; it is that the model is wrong *fluently*, in clean prose that reads exactly like the truth. IROs are trained to distrust polish that arrives too easily, and they are right to.

## Fear 2: Governance

Who approved this output? What data was fed in to produce it? Did anything non-public leave the building? Can we reconstruct, after the fact, how a number in a board deck was generated? Most consumer-grade AI use cannot answer a single one of those questions. For a function that lives and dies by Reg FD, selective-disclosure discipline, and an auditable record, an ungoverned tool is not a productivity gain; it is an unmonitored liability.

## Fear 3: The second full-time job

This is the fear no one puts on a slide, and it is the most decisive of the three. The IRO already runs a full calendar: earnings, targeting, the board, the inbound. The unspoken dread is that “adopting AI” really means *becoming the person who builds, populates, prompts, supervises, and maintains an AI system* on top of everything else. The tool stops being help and becomes a project. And projects, on a full calendar, quietly die.

### WHY THIS MATTERS

These three fears explain almost every stalled rollout in the market. A model that solves the work but ignores accuracy, governance, and the IRO's time will be abandoned, no matter how good the demo looked.

*The IRO's instinct is sound: a tool that adds risk and adds work, however clever, is not a tool. It's a liability with a nice interface.*

## SECTION 03

## Why most AI pilots fail

Failed IR pilots rarely fail loudly. They are not cancelled; they are quietly stopped being used. The pattern is consistent enough to be predicted in advance.

- 1 They start with a tool, not a workflow.** The team buys capability, “an AI platform,” rather than committing to an outcome, like “the board capital–markets report writes its own first draft every quarter.” Capability with no destination produces demos, not results.
- 2 No one owns it.** The pilot is everyone’s interesting side project and no one’s actual responsibility. The moment earnings season hits, the side project loses every time.
- 3 The data foundation isn’t there.** The system is pointed at scattered, inconsistent, out–of–date source material, and confidently reflects that mess back. Garbage in, fluent garbage out.
- 4 It lives outside the cadence.** The pilot runs beside the real IR calendar instead of inside it. If it isn’t embedded in the earnings cycle, the targeting motion, and the board rhythm, it is an experiment, and experiments are easy to skip.
- 5 It’s judged on novelty, not on a saved hour.** “Look what it can do” wins the meeting; “it reliably saved me a real afternoon, on real work, that I trust” wins the year. Most pilots are never held to the second standard.
- 6 The demo gap.** Tools that dazzle on a clean sample document turn brittle against a real company’s decade of inconsistent filings, renamed metrics, and competing narratives. The gap between the demo and the data is where pilots go to die.

**The common thread:** nearly every failure traces back to treating AI adoption as a *software decision* rather than an *operating decision*. The teams that succeed don’t buy a smarter tool. They redesign how a recurring piece of work gets done, and put someone with IR judgment in charge of it.

*A pilot that survives earnings season was never really a pilot. It was the start of an operating model.*

SECTION 04

# What top IROs automate, and never will

The best practitioners are not the most aggressive automators. They are the most *deliberate* ones. They have a clear, defensible line, and they hold it.

The principle is simple to state and hard to live: **automate the gathering, never the deciding.**

AI does the legwork: the watching, retrieving, summarizing, and drafting. The IRO does the judgment: what it means, what to say, what to disclose, and what to send. The machine proposes. The practitioner disposes.

AUTOMATE FREELY	NEVER WITHOUT A HUMAN
+ Continuous monitoring of ownership, peers, and short interest	× The final word on any external disclosure
+ Earnings-call and analyst-note summarization	× What is, and isn't, material
+ Peer and competitor filing surveillance	× Guidance, and any number that ships
+ First drafts of scripts, Q&A, and FAQs	× The judgment call on selective disclosure
+ Investor targeting lists and portfolio-fit screens	× The relationship with a specific analyst or PM
+ Meeting-prep briefs and management bios	× The strategic narrative decision
+ Sentiment and narrative-consistency checks	× Anything that touches material non-public information
+ Data aggregation for the board report	× Pressing "send"

What makes this line durable is that it maps cleanly onto the three fears. Automating the gathering attacks the *second-job* fear: the drudgery leaves. Keeping a human on every decision answers the *accuracy* and *governance* fears: nothing reaches an investor, a filing, or the board without a named person standing behind it.

THE OPERATING RULE

AI should make the IRO's judgment go further, not stand in for it. The moment a system is allowed to decide rather than to *recommend*, it has crossed the line that the best IR programs never cross.

## SECTION 05

# The future of IR operating systems

The next shift in IR is not a better chatbot. It is the move from *tools you open* to a *system that runs*: continuously, in the background, on a foundation that remembers.

Today's tools are episodic. You go to them, ask a question, get an answer, and leave. Nothing persists. An IR operating system inverts that relationship. It holds your context as a living foundation: your filings, your investor base, your peer set, your history, your narrative. It works against that foundation whether or not anyone has a tab open.

The maturity curve runs in four stages. Most of the market is still on the first two.

<p>STAGE 01</p> <h2>Assist</h2> <p>Ad-hoc chat. The IRO asks; the model answers. Useful, personal, nothing accumulates.</p>	<p>STAGE 02</p> <h2>Augment</h2> <p>A knowledge base of the company's own materials. Answers get grounded, but the IRO still drives every step.</p>	<p>STAGE 03</p> <h2>Operate</h2> <p>The system runs the recurring work continuously (monitoring, drafting, prepping) and hands the IRO decisions, not blank pages.</p>	<p>STAGE 04</p> <h2>Orchestrate</h2> <p>A coordinated layer across the full IR calendar, with the IRO directing outcomes and the system carrying the load.</p>
---	---	--	--

The decisive change at Stage 03 is a change in the IRO's posture. The job stops being *doing the work* and becomes *directing it*: reviewing, deciding, and signing off on work the system has already advanced to the edge of completion. The practitioner moves up the value chain. The system does the reaching; the human does the choosing.

*The future IRO doesn't open a tool and start typing. They open a brief that's already written, and decide what to do about it.*

## SECTION 06 · THE MODEL

# What an AI-native IR operating model looks like

Everything in this report points to one conclusion: the winning model is not a smarter tool bolted onto an old workflow. It is a different operating model, one designed, from the first line, around the way investor relations actually works.

We have spent the last several years building and running IR this way. The principles below are not aspirations. They are the design rules we hold ourselves to, and the ones we believe will define serious IR programs for the next decade.

## 01 Built around the calendar, not bolted onto it

The system is organized around the real rhythm of IR (the earnings cycle, the targeting motion, the board cadence), not around a generic chat box. It meets the work where the work already lives.

## 02 It watches so the IRO doesn't have to

Ownership shifts, peer moves, sentiment changes, and strategic triggers are monitored continuously. The IRO doesn't start from a blank page; they start from a brief, with the relevant signal already surfaced and scored.

## 03 It recommends; it never decides

Every consequential output is a recommendation for a human to approve, edit, or reject. The system drafts the outreach; it does not send it. Green-light gating isn't a setting; it's the architecture.

## 04 It never fabricates, and it shows its work

Outputs are grounded in real sources, with provenance attached. Where something is inferred rather than known, it is flagged as an inference, not laundered into false confidence. Every claim can be traced back to where it came from.

## 05 It amplifies senior judgment

The model is built to make an experienced practitioner's judgment reach further, not to replace headcount with unsupervised output. The human in the loop is the most senior person available, not the most junior.

## 06 It is governed by default

What went in, what came out, who approved it, and where the data lives are all answerable, before anyone asks. Accuracy and governance aren't features added later; they are conditions of the system existing at all.

### READ THE PRINCIPLES AGAIN

Each one is an answer to a fear from Section 02 or a failure mode from Section 03. That is not a coincidence; it is the point. An AI-native model earns its place by retiring the exact risks that kill ordinary pilots.



GlobalIRGroup

## IN CLOSING

# The work is changing. The judgment isn't.

AI will not replace the investor relations officer. But it has already changed what an IRO spends a day doing, and over the next few years, the programs that pull ahead will be the ones that hand the gathering to the machine and reserve the judgment for the practitioner.

That is the model we build, run, and stand behind. If this report resembled the way you already think about the function, you are the kind of operator we tend to work well with.

### PRIVATE DEMONSTRATIONS

We don't run a public pitch. For IR leaders and management teams who want to see an AI-native IR operating model working on real capital-markets workflows, private demonstrations are available on request.

[don@globalirgroup.com](mailto:don@globalirgroup.com)

*No deck to download. No drip campaign waiting on the other side of a form.*

*Just a conversation, IRO-to-IRO, when you're ready to have it.*