The modern compliance team playbook

For compliance and HR leaders working to close measurement gaps, run a more effective program, and get their team's time back.

Introduction

A note before you read

The Measurement Gap Assessment told you where your program stands. This playbook tells you what to do about it. If you have not completed the free Measurement Gap Assessment yet, we highly suggest that you do.

It's organized around the three measurement themes from the assessment: reach and trust, investigation and review quality, and program intelligence. Each section covers the activities that close the gap in that theme and, where relevant, the platform capabilities that make those activities operationally feasible.

Read the section most relevant to your priority gap first. Then read the others when you're ready.

Framing principle

What's safe to automate

Before getting into activities, one principle worth anchoring on.

The work that's safe to automate is structurally repeatable work that doesn't require case-specific judgment, i.e. triage criteria, routing rules, translation, renewal reminders, document structure, and report generation. These are activities where the right output can be defined in advance, applied consistently, and reviewed if needed.

The work that stays human is the work that requires interpretation. What the case means. What the right outcome is. How to engage with a sensitive reporter. What story the data tells the board. How to handle a situation the system hasn't seen before.

The practical test

If a senior compliance leader and a junior team member would arrive at the same answer given the same input, it's safe to automate. If their answers would differ in ways that matter, it isn't.

The activities in this playbook draw that line for each area of the program.

Section 1

Reach and trust

Six KPIs across two questions: is the program being heard, and do people who engage with it trust it enough to come back?

Section 1 · Reach and trust

What this theme measures

Six KPIs across two questions: is the program being heard, and do people who engage with it trust it enough to come back? Reports per 100 employees, intake method mix, time to report, anonymous-versus-named ratio, check-back rate, and disclosure campaign participation rate.

A gap in this theme usually means the program is receiving less signal than the organization is generating. The cases arriving look fine. The cases not arriving are invisible.

How to measure it

Your analytics dashboard already holds most of what you need: volume, channel mix, anonymous-versus-named ratios, check-back rates, disclosure participation. The question is how quickly you can find the patterns inside it.

Standard analytics gives you the data. Sienna Insights makes uncovering the pattern faster: ask which regions are trending below their expected reporting rate, how your check-back rate has shifted over the last four quarters, where your disclosure participation is falling short by team. Same data. Less time between the question and the answer.

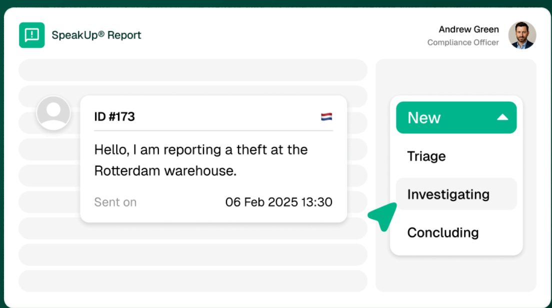

SpeakUp Report — *reporters submit through any channel and follow case status; the two-way exchange that drives check-back rates happens right here.*

Activities

01 Open more channels, including voice and mobile

A single reporting channel systematically excludes parts of the workforce. Frontline employees, non-desk workers, people whose first language doesn't match the form's language, and people in cultures where direct digital submission doesn't feel natural will underreport. The reach gap this creates is invisible until you normalize volume by headcount and look at it by population.

SpeakUp Report provides multichannel intake across web, mobile app, phone, voice agent, and manager-routed flows. The voice agent makes the channel accessible to employees who would never use a web form: deskless workers, lower-literacy populations, anyone who'd rather speak than type. Each channel is available in multiple languages at point of submission.

If you currently receive reports from one or two channels and have populations that match the profiles above, adding the relevant channel typically lifts normalized reporting volume within a quarter.

This activity impacts: reports per 100 employees (KPI 1), intake method mix (KPI 2)

---

02 Normalize reporting volume with HR data

Raw case counts are almost meaningless for comparison across regions or business units. A region with 30 reports could be your healthiest if it has 2,000 employees, or your quietest if it has 200. Normalizing by headcount turns the count into a rate, and the rate is what surfaces silent regions.

You don't need a deep system integration to do this. Most compliance teams that produce normalized reporting bring case data and HR headcount data together inside their own reporting environment, typically a Power BI dashboard or similar. SpeakUp Report can export case data into platforms like SharePoint and Power BI on a defined cadence, where your team combines it with HR headcount and location data on your side. The setup is straightforward and the resulting metric (reports per 100 employees, broken out by region) is one of the most diagnostic in the framework.

This is often the single highest-impact operational step a program with a Reach and Trust gap can take. The rest of the activities in this section become more interpretable the moment the denominator is accurate.

This activity impacts: reports per 100 employees (KPI 1), geographic distribution (KPI 15), disclosure participation rate (KPI 6)

Activities

03 Improve communication with reporters at every milestone

A reporter who hears nothing after submitting tends not to come back. And a reporter who doesn't come back tends not to report again. Low check-back rates (the industry baseline is 40%+ for anonymous reporters, 60%+ for named) are usually a communication problem more than a trust problem: the program isn't telling reporters what's happening.

The fix is operational discipline at the case level: a status update when the case is received, when it's under review, when additional information is needed, when the outcome is reached. SpeakUp Report's reporter portal supports this through structured two-way messaging, with the reporter accessing the case via a personal access code regardless of whether they reported anonymously or named. Whether updates are sent manually by handlers or built into your standard case workflow, the goal is the same: reporters never wait in silence.

Programs that move from inconsistent communication to a defined milestone cadence typically see check-back rates climb above the 40% benchmark within one to two quarters.

This activity impacts: check-back rate (KPI 5), anonymous-versus-named ratio (KPI 4) over time

---

04 Run structured proactive disclosure campaigns

Without formal disclosure campaigns, your program has no proactive view of conflicts of interest, gifts, outside employment, or related-party relationships. The risks these would have surfaced still exist; they're just invisible until something goes wrong.

SpeakUp Paths runs structured disclosure campaigns with reminder cadences, manager visibility into team completion, and mobile-accessible submission for distributed workforces. The Employee Hub gives employees a single place to submit, track, and respond to their disclosures without routing through compliance. Participation rates, renewal completion, and mitigation follow-through are visible in the analytics dashboard throughout the campaign.

Programs moving from informal annual campaigns to structured Paths campaigns typically see participation rates climb by 15 to 20 percentage points within two cycles. More importantly, they generate a program-health KPI they previously couldn't produce at all.

This activity impacts: disclosure participation rate (KPI 6), disclosure renewal and follow-through (KPI 13)

Section 1 · Outcomes

Closing this gap

A program that has worked through the activities in this section will recognize itself in the outcomes below. Treat them as a target state, not a checklist to race through.

What good looks like

- Reporting rate normalized by headcount and tracked quarterly, with no region sitting materially below its expected rate without a credible explanation
- Intake available across three or more channels, with uptake from populations that previously underreported
- Check-back rate above 40% for anonymous reporters and 60% for named, sustained across case types
- Annual or quarterly disclosure campaigns running on a defined cadence, with manager visibility into team completion and participation above 85%

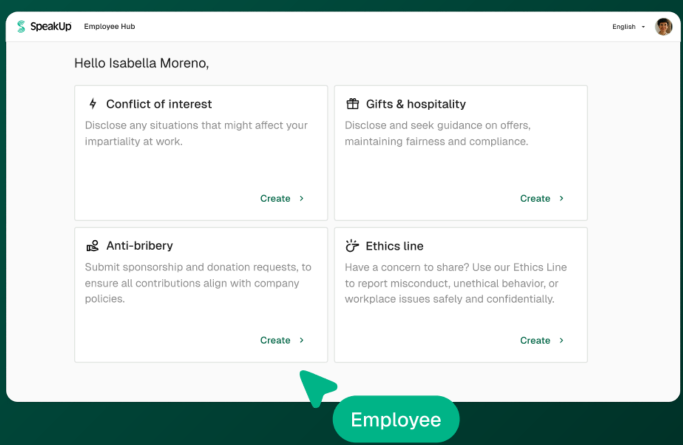

SpeakUp Paths — *the Employee Hub gives staff one place to file conflict-of-interest, gifts, and anti-bribery disclosures, on a defined campaign cadence.*

Section 2

Investigation and review quality

Seven KPIs across reactive case handling and proactive disclosure review.

What this theme measures

Seven KPIs across reactive case handling and proactive disclosure review: average time to reply, repeat reporter message rate, case closure time, substantiation rate, time to disclosure decision, disclosure outcome distribution, and disclosure renewal and mitigation follow-through.

What ties them together is operational consistency. The work isn't difficult in the abstract. It's difficult to do at the same standard, across handlers, across case types, across regions, week after week.

How to measure it

The data for these KPIs is already in your case and disclosure records. Sienna Insights makes it interrogable without an analyst: ask how your substantiation rate is distributed across handlers, where the bottlenecks are in your case workflow, how closure time has moved over the last year, or how your disclosure outcome distribution compares across regions. Patterns that previously showed up in quarterly reviews start surfacing weekly.

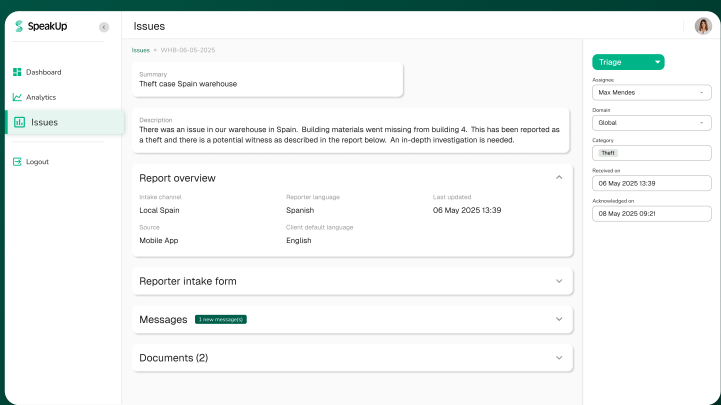

SpeakUp Report — *triage status, assignee, category and the structured report overview for a single case, in one workflow view.*

Activities

01 Build a baseline case workflow with configurable issue types

The single highest-leverage operational change in this theme. Programs at Level 2 typically have workflows that exist on paper but vary in practice: different handlers, different documentation standards, different decision logic. Closing the gap means defining a workflow every case follows, with required fields at each stage, response SLAs, and escalation rules.

SpeakUp Report lets compliance teams configure workflows themselves: stage definitions, required fields, routing rules, SLA thresholds. No IT involvement. Issue types are configurable to match your actual risk landscape rather than the default taxonomy, which means your category data reflects what your organization actually sees.

Without a defined workflow, automation has nothing to enforce.

This activity impacts: case closure time (KPI 9), substantiation rate (KPI 10), average time to reply (KPI 7)

---

02 Automate triage and routing

First response time is largely a triage problem. Cases sit in a queue because triage requires reading, categorizing, severity-rating, and routing. None of that requires a compliance expert once the criteria are established.

Sienna's AI triage reads each incoming case at submission and produces a structured summary, suggested category, severity rating, and routing based on the rules your team defines. Handlers receive cases ready to act on rather than ready to assess. The triage step compresses from 15 to 30 minutes per case to under two minutes of review.

If your team currently spends ten hours a week on triage, that typically compresses to two to three hours. The saved time goes directly into case review and handler judgment.

This activity impacts: average time to reply (KPI 7), case closure time (KPI 9), handler workload distribution

Activities

03 Reduce the language barrier in cross-language cases

Multinational programs almost always carry a quiet drag on case quality from cross-language handling. Reporters in their non-native language produce thinner accounts. Handlers communicating in the reporter's non-native language miss nuance. Cases sit longer, message counts spike, reporters disengage. This rarely shows up on any report until it shows up in closure time and message rate trends.

Translation that operates inside the case workflow rather than as an external service is what removes the friction. SpeakUp Report today supports multilingual reporting on the reporter's side: reporters submit in their preferred language and engage with the platform in that language at point of submission. Bidirectional in-case translation between reporter and handler is on the roadmap and expected later this year, which will let handler messages translate automatically the other direction without external coordination.

Even with the current capability, the operational effect is meaningful: reporters engage more openly when they aren't translating their experience into a second language to submit it, and message rates compress as a result. Once bidirectional translation is live, the back-and-forth between handler and reporter on cross-language cases will run at the same speed and quality as native-language cases.

This activity impacts: repeat reporter message rate (KPI 8), case closure time (KPI 9), check-back rate (KPI 5)

---

04 Use AI summarization to accelerate case review

Reading a full case file before review or escalation takes time that compounds across a busy quarter. If a typical review requires twenty minutes of reading and you're handling fifty cases per quarter, that's sixteen hours before a single decision is made.

Sienna AI summarization gives reviewers a structured brief: what happened, who's involved, relevant context, what's been done so far. Reviewers start from a base rather than a blank page — typically twelve to thirteen hours back per quarter on the fifty-case example.

Anonymization is embedded in the same capability. Identifying details are stripped automatically when a case is summarized for cross-team review, which is what makes handler calibration operationally safe.

This activity impacts: review and escalation efficiency, case closure time (KPI 9), substantiation rate calibration (KPI 10)

Activities

05 Capture decision rationale as required fields

Outcome without rationale is a number you can't defend. Capturing outcome, rationale, reviewer, and date as required fields on every case and disclosure decision is what makes substantiation rate defensible to a board, disclosure outcome distribution interrogable, and handler calibration possible. It's also what makes it impossible to close a case or a disclosure without documenting the decision, which is the enforcement mechanism that turns the discipline from optional to consistent.

SpeakUp Report and SpeakUp Paths both treat these as structured required fields at the decision point, not as free-text notes. Records are retained and retrievable on demand. When an auditor or regulator asks why a specific decision was made, the answer is in the system in seconds rather than reconstructed over days.

This activity impacts: substantiation rate (KPI 10), disclosure outcome distribution (KPI 12), audit readiness

---

06 Track renewals and mitigation actions to completion

A program that runs good disclosure campaigns but doesn't track renewals or verify mitigation completion is running campaigns that don't hold. Yesterday's mitigated conflict of interest may already be unmanaged risk if nobody confirmed the mitigation actually happened. This is the gap between a participation rate that looks good and a program that's actually managing ongoing risk.

Closing this gap is operational. It means defining renewal cadences against disclosure dates, assigning ownership for verifying that mitigation actions actually happened, and building a check-in rhythm with managers whose teams have active disclosures. Without that discipline, the data from a successful campaign decays inside two quarters.

If your team currently spends five to ten hours a week chasing renewals manually, putting structure around the process won't eliminate that time entirely, but it will make it productive: time spent on the disclosures that actually need attention rather than on tracking down who has and hasn't responded. Completion rates typically lift from a 60 to 70% range to above 90% within two renewal cycles when ownership and cadence are in place.

This activity impacts: disclosure renewal and mitigation follow-through (KPI 13)

Section 2 · Outcomes

Closing this gap

A program that has worked through the activities in this section will recognize itself in the outcomes below.

What good looks like

- Case closure time at or below the 30-day industry benchmark, with bottleneck visibility for cases aging past threshold
- Substantiation rate stable in the 30 to 60% range, defensible to auditors, comparable across handlers with rationale captured on every decision
- First response time consistently within SLA, with most pre-triage time redirected to case work
- Cross-language cases running at improving quality and speed, with the gap to native-language cases narrowing as translation capabilities mature
- Disclosure outcome distribution showing meaningful variation across approved, mitigated, and prohibited
- Renewal and mitigation follow-through above 90%, with active disclosures connected to ongoing case management

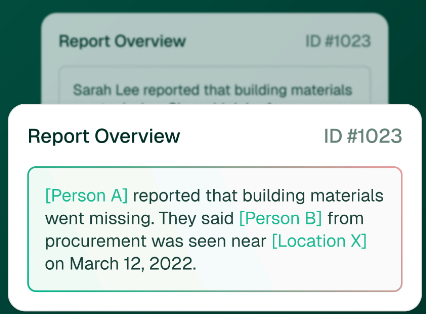

Sienna AI — *identifying details are stripped automatically when a case is summarized for cross-team review and handler calibration.*

Section 3

Program intelligence

Two KPIs that tell you what your compliance data is revealing about where risk is moving.

What this theme measures

Two KPIs: issue category distribution and trend, and geographic and entity-level distribution. Together they answer a single question: what is your compliance data telling you about where risk is moving before it surfaces somewhere else?

A gap in this theme doesn't mean the data isn't there. It almost always means the data is there but the program isn't yet organized to ask the right questions of it.

How to measure it

This theme is where Sienna Insights does its most distinctive work. Instead of pulling data into a spreadsheet and looking for patterns, your team asks questions in plain language and gets answers in seconds. "Which categories have increased most in the last two quarters?" "Which regions are reporting below their expected rate given headcount?" "How has the mix of issue types shifted since we expanded into that new market?" The analysis that previously took a day now takes a conversation.

Activities

01 Configure issue types to match your actual risk landscape

Most programs start with the category taxonomy the platform shipped with. That taxonomy is generic by design. A manufacturing program with supply chain exposure needs different sub-categories than a financial services program focused on conduct risk. Using a generic taxonomy produces data that's hard to act on.

Configuring issue types in SpeakUp Report takes minutes rather than a development cycle. Your team sets the categories, sub-categories, and severity levels that reflect your organization's actual risk landscape. Once the taxonomy is right, category distribution and trend data (KPI 14) becomes interpretable: patterns surface in the right buckets and early-warning signals are where you'd expect them.

Refining the taxonomy against the last 12 to 24 months of case data is the starting point. Most programs find three to five categories that need splitting and two to three that should merge.

This activity impacts: issue category distribution and trend (KPI 14)

---

02 Interrogate patterns with natural-language questions

Quarterly category reviews are the standard practice. Most programs produce a table of case counts by category and look for spikes. What they miss is the drift: a slow, sustained increase in a particular category over three or four quarters that doesn't show up as a spike in any single quarter but is clearly directional when you look across time.

Sienna Insights surfaces this without requiring a structured data analysis request. Ask the question in plain language and the answer comes back with the relevant data. "Which categories have moved most over the last four quarters?" "Where is procurement-related reporting trending?" "Is the increase in manager conduct cases concentrated in a particular region or spread across the organization?" The questions your team would previously save for an analyst (or skip because there wasn't one) become part of a weekly review.

The shift this produces isn't just efficiency. It's timing: patterns that previously surfaced in incident reports now surface in your compliance data months earlier.

This activity impacts: issue category distribution and trend (KPI 14), geographic distribution (KPI 15)

Activities

03 Build a review rhythm that catches drift before it spikes

Quarterly reviews catch patterns in retrospect. The point of program intelligence is catching them as they develop, which means the cadence of how often you're looking matters as much as the depth of any single review.

Most programs that close this gap establish a lighter monthly rhythm alongside the quarterly review: ten or fifteen minutes asking the same questions about category and regional movement. The questions stay consistent so a real change shows up against a steady baseline.

Define what counts as meaningful: a sustained increase in a category over two consecutive months, a region whose reporting rate has dropped meaningfully, a spike in a sensitive category. The thresholds are yours to set; the discipline of looking on a defined cadence is what turns the data into early warning.

This activity impacts: issue category distribution and trend (KPI 14), geographic distribution (KPI 15)

---

04 Build geographic analysis on normalized data

Raw counts by location make geographic comparison meaningless. A region with 30 reports could be your most active or your quietest depending on headcount. The analysis only becomes useful when normalized by headcount.

The setup mirrors Activity 2 in Section 1. SpeakUp exports case data into your reporting environment (Power BI or similar), where your team combines it with HR headcount and location data. Sienna Insights can then cross-cut by region and category: which issue types are most common where, which regions are silent relative to headcount. Silent regions are almost never a clean culture — they're almost always a trust or accessibility problem worth investigating.

This activity impacts: geographic and entity-level distribution (KPI 15), reports per 100 employees (KPI 1)

Closing this gap

A program that has worked through the activities in this section will recognize itself in the outcomes below.

What good looks like

- Issue category taxonomy configured to reflect your actual risk landscape, with sub-categories that make pattern analysis meaningful
- Category trends reviewed on a regular cadence, with drift visible across quarters rather than only at spikes
- A monthly review rhythm with defined thresholds for what counts as meaningful movement, sitting alongside the deeper quarterly review
- Geographic distribution normalized by headcount in your reporting environment, with no region materially below its expected rate without explanation
- The compliance team using Sienna to ask questions of the data in real time, not waiting for a quarterly report to tell them what's in it

Where to start

The compounding effect

Closing one gap accelerates closing the others. The connection is structural.

When reach and trust improves, more signal enters the program: more reports, more disclosures, better representation across populations. That signal is the input to investigation and review quality, which means the second theme gets harder to manage without the first being solid. And when investigation and review quality improves, the data those cases generate becomes richer and more consistent, which is what makes program intelligence finally worth interrogating.

The same logic runs through how the platform supports the work. Configurable issue types and workflows in Case Management produce the structured data that Sienna can interrogate. Sienna's triage, translation, and summarization capabilities work because they're built on top of the same case and disclosure records that SpeakUp Report and Paths create. Each part of the platform makes the others more useful.

Where to start tomorrow

Follow your priority gap from the assessment. The section most relevant to your lowest score is where the fastest improvement is available, and where the operational work and platform capabilities produce the most immediate time recovery.

If you had a gap across multiple themes, start with reach and trust. It's the input to every other theme: a program that isn't being heard can't produce reliable program intelligence, and the cases that aren't arriving can't be investigated with rigor.

Within whichever theme you start with, the activities that produce the fastest improvement are almost always the ones that compress a manual, repeatable step: triage, translation, renewal cadences, taxonomy refinement. Start there.

The platform supports continuous improvement

The goal isn't to run a project that closes a gap. It's to build a program where measurement and operational work reinforce each other continuously: the data tells you where the program needs attention, and the platform makes it feasible to give it that attention.

SpeakUp is built for that. One platform across intake, case management, proactive disclosure, and AI-powered analysis.

Get started

Ready to see this in your program?

Book a guided demo framed around your priority gap. We'll show you how the SpeakUp Platform produces the KPIs relevant to your specific measurement gap and what your team gets back when it does.

Or if you'd prefer to explore the SpeakUp Platform at your own pace and with no strings attached, try our self-guided interactive demo.

Book a guided demo →

Try the self-guided interactive demo →