RingCentral

US

Enterprise AI governance in communications

Why it matters and how to get ready



Executive overview



Artificial intelligence is now embedded in everyday business communications: meetings are summarized automatically, sales calls are transcribed and analyzed, and customer service agents receive suggested responses in real time. These advances promise efficiency, better experiences, and faster decisions. Yet they also raise new questions:

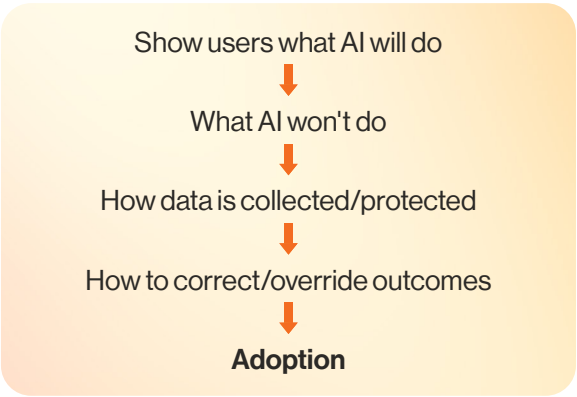
- What data is being used?
- Who is accountable if an AI system makes a mistake?
- How do we demonstrate that our use of AI is fair, secure, and compliant across jurisdictions?

AI governance is the discipline that provides credible answers to those questions. It defines how your organization designs, deploys, and measures AI so that benefits are realized without compromising safety, privacy, ethics, or compliance.

For CIOs, CTOs, and other enterprise decision-makers, governance is not a bolt-on. It is a business capability that builds trust with customers, employees, investors, and regulators. This paper explains why governance matters, outlines practical steps to get ready, and describes RingCentral's approach to responsible AI for enterprise communications.

Why AI governance matters now

Trust is the adoption bottleneck



Employees and customers don't adopt what they don't understand or can't control. Employees and customers will embrace AI in communications when the experience makes its boundaries unmistakable: what the system is designed to do, what it will never do, how their data is collected and protected, and how they can correct or override outcomes. That clarity isn't a policy PDF. it's built into the product through plain-language disclosures, timely consent prompts, visible "AI-generated" labels, simple review-and-edit flows, and reversible actions for anything consequential. It's reinforced for administrators with a transparency center that shows which features are on, how they're configured, and how they're performing. And it's sustained by measuring trust signals such as opt-in rates, correction rates, and time to resolve AI-related issues. When governance shows up this way in the experience, confidence compounds and adoption grow. When it doesn't, even great features stall under skepticism, shadow rules, and avoidable escalations.

Regulations are converging

Across markets, policymakers are introducing risk-based rules for transparency, data protection, human oversight, and post-market monitoring. Whether you look to the EU's comprehensive approach or sector rules in the U.S. and elsewhere, the themes are consistent: be transparent about when and how AI is used, protect the data you touch, keep humans in the loop for consequential decisions, and watch systems after launch with the same rigor you applied before. Dates and definitions may shift, but the expectation won't: you must show your work. That means maintaining clear documentation of use cases and data flows, publishing plain-language user disclosures, keeping evidence of evaluations and approvals, recording changes to prompts and models, and having a tested plan to detect issues, explain what happened, and correct them quickly. Enterprises that build this muscle now can satisfy multiple regulators with one set of artifacts, reduce audit friction, and move faster precisely because their governance is durable and portable.

Communications data is uniquely sensitive

Voice, video, and messages capture not only the words exchanged but also who said them, to whom, when, and in what tone. Voice recordings can reveal identity traits and health or biometric indicators; video can expose surroundings and bystanders; chat threads often blend customer PII with trade secrets and decisions in flight. That richness makes the data powerful for AI, and fragile. A mis-summarized meeting, an over-shared transcript, or an unvetted model integration can leak intent, pricing, or personal details in seconds. Unlike static documents, calls and messages happen in real time, shrinking the window to prevent or remediate harm. Robust guardrails are therefore nonnegotiable: clear, jurisdiction-aware consent; least-privilege access; default minimization and redaction; purposeful retention with legal holds; and visible provenance for AI-generated artifacts. When these controls are absent or inconsistent, the blast radius is large and immediate: customers lose trust, negotiations are compromised, and regulatory exposure accelerates.

The cost of rework is high

Invest early to save cycles later

Retrofit later	Build early
Slow	Fast
Costly	Reusable
Erodes trust	Preserves momentum

Clear guardrails means faster launches

Every missing control becomes expensive technical and organizational debt. When consent prompts, disclosures, retention rules, or approval gates are bolted on after launch, engineers must refactor code and data flows, product teams pause roadmaps, legal rewrites user communications, and support absorbs avoidable confusion. Users experience shifting behaviors—labels appear, features change, access tightens—eroding trust just as you're trying to scale adoption. Audits become slower and more adversarial because evidence is scattered and retroactive. By contrast, a light but explicit governance foundation—standard consent and disclosure patterns, clear autonomy tiers, predefined evaluation and sign-off steps, and an admin transparency center—turns compliance into a repeatable product ritual. Teams ship faster because expectations are known up front; artifacts are reusable across features and reviews; and course corrections are smaller, cheaper, and less visible to end users. In practice, investing early saves cycles later and preserves the momentum that makes AI programs successful.

What decision-makers need to prepare

Establish ownership and accountability

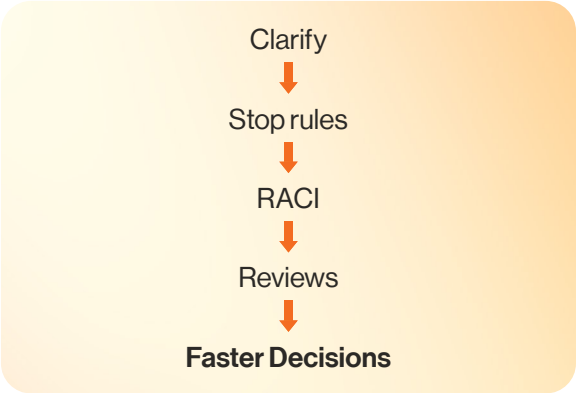


Governance works when someone is clearly in charge. Start by making AI a line of business accountability, not a side project. Assign a senior leader to own outcomes, budget, and cadence, and make it explicit that product, security, privacy, legal, and operations share responsibility for how AI is designed, launched, and managed. Publish a simple charter that states the purpose of governance, the scope of AI use it covers, and the cadence of reviews. Then make it visible: a named point of contact, a standing forum, and a lightweight reporting rhythm to the C-suite and board. When people know who owns decisions and how progress is measured, they move faster and escalate less. The goal is not bureaucracy; it's creating a predictable path from idea to safe deployment so teams can innovate with confidence.

Name an executive sponsor

Name an executive sponsor and create a cross-functional AI Governance Council (product, security, privacy/legal, compliance, data, and business units). Put one accountable executive at the helm—CIO, CTO, CISO, or a business unit leader—who can reconcile speed, risk, and resources. Surround that sponsor with a cross-functional council drawn from product, security, privacy/legal, compliance, data, and key business units. This council should approve use cases, set risk thresholds, and review evidence before launch and at regular intervals after. Keep meetings short and focused: a shared dashboard, a queue of proposals, and clear go/hold decisions. Treat the council as a product accelerator, not a gatekeeper—helping teams shape use cases to meet standards quickly. Publish decisions and rationales so everyone learns the pattern. Over time, this muscle memory reduces friction: teams arrive with better proposals, reviewers spot issues earlier, and leadership has consistent visibility into benefits and risks.

Define decision rights



Ambiguity slows launches. Clarify who signs off at each stage of an AI feature's lifecycle—intake, data use, design, testing, launch, monitoring, and retirement—so there are no last-minute surprises. Make those decision rights explicit and visible: product owns business fit; security and privacy own data treatment and controls; legal owns disclosures and terms; operations own rollout and incident response. Agree on "stop the line" conditions and who can invoke them. Pair this with a simple RACI for routine changes versus material ones, so minor iterations don't trigger full re-approval. Finally, commit to post-launch reviews at set intervals, with the authority to tighten controls, pause features, or retire them. When decision rights are clear, risks are addressed early, approvals are faster, and everyone understands how to course-correct if the real world disagrees with the plan.

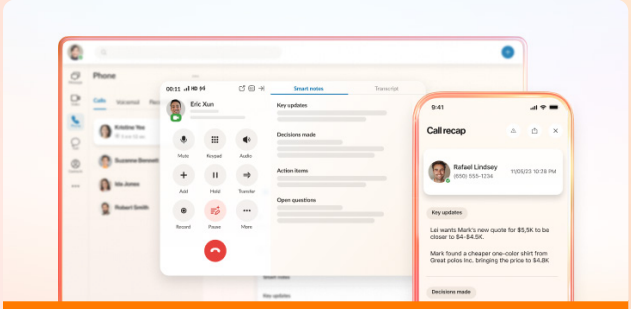
Set clear principles and translate them into policies

Principles set the tone; policies set the rules. Start with a short, memorable set—lawfulness, accountability, transparent by default, privacy and security by design, robustness, fairness, and continuous improvement. Then turn each principle into operational policies people can apply without a lawyer at their elbow. Specify which data can be used and for what purpose; which use cases are allowed, restricted, or prohibited; what disclosures and consent are required in each channel; what level of human oversight is necessary based on risk; and how incidents are identified, escalated, and communicated. Keep policies practical: include examples, default settings, and decision trees. Bake them into templates and tools—standard consent prompts, labeling patterns, approval checklists—so teams follow the rules by default. The outcome is a shared playbook that makes expectations predictable and repeatable, speeding delivery while strengthening trust.

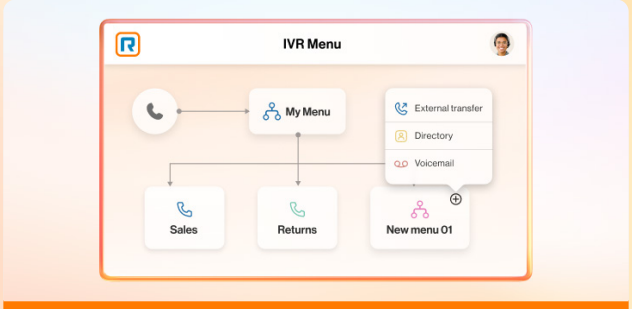
How RingCentral makes responsible AI the default (what this means in practice)

If you remember one thing, let it be this: governance should show up in the product, not just in a policy. Our approach is to make responsible AI feel obvious to users and configurable for admins—so trust grows with use, not with slideware. Practically, that starts with consent and disclosure woven into everyday moments. When transcription or summaries are on, people see it, understand why, and can opt out. When AI drafts a note or a reply, it's labeled and easy to edit—because the point is to help, not to hide. On the admin side, we give you a clear "cockpit" to set your guardrails without code: choose retention (30/90/365 days), turn redaction on by default, set your region preferences, and—where choice exists—pick your model provider without losing protections. Your data stays inside tenant boundaries; we don't train models on your content unless you explicitly opt in; and any change to these settings is logged so audits feel like a review, not an expedition.

Governance should show up in the product, not just in a policy.



For Users:
Transparency, labels, reversibility



For Admins:
Control, guardrails, auditability

We also keep control human and right-sized. Think in simple autonomy tiers—suggest, draft, act-with-confirm, act-with-rollback—and match them to the consequence of the action. Most features sit at "suggest" out of the gate; anything that touches customers or regulated data asks for a human "yes" before it acts; and low-risk steps get a safety net with one-click undo. This isn't red tape—it's speed with a seatbelt. To know it's working, we measure a handful of signals that leaders care about: time saved, adoption and satisfaction, user-rated accuracy of summaries, acceptance of suggested replies, consent coverage, and time to resolve issues when they pop up. If a metric dips, we don't debate philosophy—we tighten filters, lower autonomy, or roll back a change and tell you what we did. Finally, because AI is a team sport, we're transparent about our vendor stance: where data is processed, what training defaults are, and how to switch off or swap a model without weakening your controls. In short, responsible by default, explainable on demand, and adjustable in minutes.



Measuring signals that matter

- Time saved
- Adoption and satisfaction
- User-rated accuracy of summaries
- Acceptance of suggested replies
- Consent coverage
- Time to resolve issues

A simple 90-day plan to get ready (so you can move fast with confidence)

Phase 1: Days 0–30
Build the foundation

Start small:
Council + principles + 2 visible use cases.

You don't need a giant program to start—you need momentum and a few visible wins. In the first month, stand up a small governance council (one accountable exec plus product, security, privacy/legal, and operations), publish five to seven plain-language principles, and inventory where AI already shows up in your communications stack. Turn on user-facing labels and consent prompts for anything live today. Pick two priority use cases—one for meetings (summaries and action items), one for the contact center (agent assist)—so you can prove value while you build muscle.

Phase 2: Days 31–60
Make guardrails real

Turn guardrails into defaults.

In days 31–60, make your guardrails real. Set retention and redaction defaults, confirm regional settings, and light up your admin "cockpit" so leaders can see what's on and how it's performing. Decide on four to six metrics per use case—value (time saved), quality (user-rated accuracy), safety (privacy flags), and readiness (incident response time). Socialize these measures with the teams who will live with them; you want feedback now, not after rollout.

Phase 3: Days 61–90
Pilot with confidence

By day 90, governance is live, measurable, and scalable.

In the final 30 days, pilot with autonomy you can trust: keep external actions at "suggest" or "act-with-confirm," and enable "act-with-rollback" only for low-risk steps. Run a one-hour tabletop exercise on an AI-related incident (a bad summary, a mislabeled consent state) to pressure-test who does what, when. Then publish a short governance report to leadership: what's live, what guardrails are in place, how it's performing against your metrics, and what you'll adjust next. By day 90, you'll have responsible AI working in production, a repeatable way to approve new uses, and evidence that speeds audits—in other words, a foundation you can scale without slowing down.

Conclusion and takeaways

Governance is how you turn promising AI features into durable business capability. In enterprise communications—where the data is personal, fast-moving, and high-impact—good governance must be visible to users, configurable by administrators, and credible to auditors. The winning pattern is straightforward: build consent and disclosure into the experience, keep data within clear boundaries, match autonomy to consequence, and measure a short list of outcomes that leaders use to steer. With this foundation, you move faster because expectations are predictable, audits are lighter, and trust compounds with every interaction. RingCentral's approach is designed for exactly this: responsible by default, explainable on demand, and adjustable in minutes—so you can scale AI with confidence across meetings, messages, and the contact center.

Takeaways

- **Govern to accelerate:** clear roles, simple policies, and reusable evidence make launches faster and safer.
- **Put governance in the product:** disclosures, consent, labels, reversibility, and admin controls build trust where it matters.
- **Calibrate autonomy:** start at suggest/draft; require confirm for consequential actions; offer one-click rollback for low-risk steps.
- **Control the data lifecycle:** collect only what you need; set retention and redaction by default; keep tenant boundaries firm.

