Your HR copilot knows who's on medical leave. Who's it allowed to tell?

$ compile --target=hr-case-mgmt --check=rbac,data-classification,adafmla

THE CASE

Claude is genuinely good at understanding what someone's really asking. The question is what happens the moment "answer it" means **disclosing something the law says only certain people may know** — a medical reason, an accommodation, a disciplinary record. That's where "Claude wrote a helpful HR bot" and "a runtime enforces who's allowed to know what" stop being the same claim.

Cost
HR case-management backend on a cloudlet: **~80% fewer AI tokens** than hand-built disclosure logic.

Quality
A disclosure the agent shouldn't make fails at the **runtime boundary**, not as a discrimination charge months later.

Security
What each role may see is enforced by the **runtime itself** — not by prompt text a reasonable-sounding question can talk it out of.

THE WRONG ANSWER IS ASKED FOR, NICELY

FROM: manager@company.com (Slack DM) — RE: Sarah Chen — coverage planning

Hey, can you catch me up on Sarah's case? **Is this a medical thing? When's she expected back — need to plan coverage.**

```
verdict: false — medical reason / diagnosis exceeds this manager's capability grant → refused, HR role required. The check doesn't grade how reasonable the question sounds — only what the role allows.
```

Claude alone: defense lives in the prompt and your guardrail code.
Magic Cloudlet: defense lives in the runtime — a well-meant question can't argue with it.

LEGEND HANDLED PARTIAL MANUAL

WHAT IT TAKES TO GET THIS TO PRODUCTION — WHO OWNS EACH STEP

	REQUIREMENT	CLAUDE ALONE, FREE-FORM	MAGIC CLOUDLET + CLAUDE
1	**HRIS connection**	MANUAL Custom integration per platform (Workday, BambooHR, ADP...).	HANDLED Generated endpoints wrapping your existing system.
2	**New-hire provisioning workflow**	MANUAL Built and maintained by your engineering team.	HANDLED Generated as endpoints — accounts, equipment, paperwork.
3	**Data classification (medical / disciplinary / comp)**	MANUAL Field-level tagging built and maintained by hand.	HANDLED A capability grant per data class, enforced at every call.
4	**Role tiers (manager / HR / HRBP / legal)**	MANUAL Hand-rolled middleware, one implementation per build.	HANDLED Runtime-native roles — same object, token to admin screen.
5	**Leave & accommodation case routing**	MANUAL Custom workflow logic for FMLA/ADA processes.	HANDLED Generated as an endpoint the correct role can call, not bypass.
6	**Over-disclosure defense**	MANUAL Filters and hardening — a defense each new question tests.	HANDLED A reasonable-sounding question still can't exceed the grant.
7	**Employee self-service scoping**	MANUAL Built by hand — what an employee sees about their own case.	HANDLED A distinct role grant, same enforcement mechanism.
8	**Audit trail of case access**	MANUAL Designed, instrumented, and wired to a log store.	HANDLED Built-in, filterable log — every access, every disclosure.
9	**Onboarding document generation**	MANUAL Template logic hand-built per document type.	HANDLED Generated from the same validated employee data.
10	**Security review before launch**	MANUAL Full human review of every generated line.	PARTIAL Structural checks handle validity; case-routing logic reviewed.
11	**Compliance scope (ADA, FMLA, GINA, EEOC, state law)**	MANUAL Ongoing, and specific to your jurisdictions.	PARTIAL Self-host/air-gap keeps infra in your boundary.
12	**Hosting, deploy, rollback**	MANUAL A DevOps project, built once a backend exists.	HANDLED Included with a managed cloudlet, or self-host / air-gapped.
13	**Ongoing patching & maintenance**	MANUAL Your team, indefinitely, as each HRIS API changes.	PARTIAL Maintained by Nuity (managed), or self-run under MIT.

13/13 manual when Claude builds the HR copilot alone. 9/13 handled and 4/13 reduced to partial when a Magic Cloudlet builds it alongside Claude.

The security model, the numbers, and the ask

what stops an over-disclosure — and what it costs either way

Five security layers, side by side

LAYER	CLAUDE ALONE, FREE-FORM	MAGIC CLOUDLET + CLAUDE
Field vocabulary	Open — any field your integration code can read.	Closed — a finite, classified set of fields per role.
Pre-execution check	Prompt hardening and app-level checks, written by hand.	Static proof vs. the live capability registry.
Capability grants	Not native — hand-built into case-summary functions.	Argument-bound (data class), checked at every dispatch.
Identity & role	Middleware each new build may implement differently.	Runtime object — one role, caller's token to admin screen.
Audit trail	Whatever your team designs and maintains.	Built-in, severity-tagged log of every disclosure.

WHERE WE'RE HONEST ABOUT TRADE-OFFS

Claude alone is faster to a rough internal prototype — a drafting assistant for offer letters or policy summaries, with no read access to case records at all. What doesn't change: a data-classification rule enforced in a prompt is a suggestion; one enforced by a runtime evaluator is a fact about what can be disclosed, regardless of how reasonable the question sounded.

What this costs, and what getting it wrong costs

	CLAUDE ALONE + YOUR TEAM	MAGIC CLOUDLET + CLAUDE
Platform cost	No fee — cost shows up in eng. time & upkeep.	Self-host free (MIT). Managed: flat $100/mo.
Build cost, ref. integration	~140K tokens to build & harden disclosure logic.	~25K tokens — ~80% fewer, generated & validated.
Time to a governed copilot	Days–weeks of review, after the demo ships.	Minutes to scoped, role-gated endpoints.

$660M
recovered by the EEOC for 17,680 workers in FY2025

$40K
average out-of-court discrimination settlement, per the EEOC

2nd
most common FY2025 lawsuit basis: disability discrimination

Source: U.S. EEOC, Fiscal Year 2025 enforcement data — agency-wide totals, not a Nuity AI estimate.

FOR THE APPROVAL MEMO

$100/month buys a hosted, role-gated HR runtime with enforced data classification and a full audit trail — a rounding error next to the $40,000 average cost of even one out-of-court discrimination settlement, before legal fees.

Twenty minutes. Bring your real HRIS.

Nuity AI's engineers will scope an onboarding and case-management copilot against your actual HRIS and role structure — live, no sales deck.

SELF-HOSTED
Free

MANAGED CLOUDLET
From $100/mo

MANAGED SERVICES
From $3,000/mo

Nuity AI
Dallas–Fort Worth, TX
clayton.redmon@nuity.ai · nuity.ai