NUITY AI

MAGIC CLOUDLETS · BUILT ON HYPERLAMBDA

BUDGET APPROVAL BRIEF

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

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

Use case 03 — HR onboarding and employee case management. A side-by-side case for pairing Magic Cloudlets with Claude, built for the person who has to approve the invoice, not just admire the demo.

- Claude alone vs Magic Cloudlet + Claude
- Why a well-meaning question is still the wrong disclosure
- What a discrimination charge actually costs

Prepared by Nuity AI

For the stakeholder taking this to procurement, security review, or finance.

nuity.ai · Dallas-Fort Worth, TX

August 2026

00 EXECUTIVE SUMMARY

An agent that answers questions about employees is a demo. One that knows what it's allowed to say is a compliance program.

Every HR leader wants the same thing: an AI that handles onboarding paperwork, tracks leave cases, and answers manager questions — without a generalist buried in tickets all day. Claude is genuinely good at that. The question this brief answers is what happens the moment "answer the question" means **disclosing something the law says only certain people may know** — a medical reason, a disability 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.

Three line items your board already tracks

Cost	Fewer engineer-hours per HRIS integration. Generating and validating the case-management backend on a Magic Cloudlet runs roughly **80% fewer AI tokens** than hand-building the same integration and disclosure logic from scratch.
Quality	A disclosure the agent shouldn't make fails at the runtime boundary, not as a discrimination charge three months later.
Security	What a manager, an HR generalist, or an employee may see about a case is enforced by the **runtime itself** — not by prompt instructions a reasonable-sounding question can talk it out of.

WHAT THIS BRIEF COVERS

Why an HR case is a uniquely hostile input for an AI agent — not because anyone's attacking it, but because the correct answer depends entirely on who's asking.

A step-by-step comparison of what it takes to get an onboarding and case-management copilot to production, Claude alone versus Magic Cloudlet + Claude.

What getting this wrong actually costs — sourced to the EEOC's own FY2025 enforcement data, not a Nuity estimate.

LEGEND — HANDLED generated & validated by the runtime PARTIAL platform assists, review still required

MANUAL your team builds and owns this

01 WHY HR IS A DIFFERENT KIND OF RISK

The wrong answer isn't hacked out of the agent — it's asked for, nicely

A support ticket is occasionally hostile. A sales negotiation is persuasive by design. An HR case is different again: the person asking usually has a completely legitimate reason to ask — and the law still says the agent may not tell them everything it knows. Nobody has to trick an HR copilot into an ADA or FMLA violation. A manager just has to ask a normal management question.

FROM: manager@company.com (Slack DM to HR copilot) 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 — I need to know if I should plan for someone to cover her role long-term or if this is short-term.** Thanks!

Neither approach assumes the manager has bad intent. The real question is what stands between that Slack message and the diagnosis actually being disclosed.

× Claude alone, free-form integration

The case-summary logic pulls from whatever fields exist on the record and paraphrases them for whoever asks. Nothing in a prompt reliably distinguishes "a manager's legitimate need to plan coverage" from "protected medical information the ADA and FMLA say only HR and the employee are entitled to see."

✓ Magic Cloudlet + Claude

Claude can draft a helpful-sounding answer. The cloudlet checks the manager's role against the case's data classification — medical and accommodation details aren't in a manager's capability grant. The response is scoped to what a manager may know: approved leave through a date, coverage needed.

```
verdict: true      "leave status & expected return date for direct report" → within manager grant → returned
verdict: scoped    "coverage duration, no medical detail" → within manager grant → returned, logged
verdict: false     "medical reason / diagnosis / accommodation detail" → outside manager's capability grant → refused, HR role required
```

Two ways to answer: what stops the agent from over-disclosing?

Ask any HR-AI vendor this question and the marketing sounds identical — "role-based," "privacy-aware," "compliant by design." The architecture underneath does not. There are, in practice, exactly two answers.

Claude alone, free-form integration

```
Question → Claude
    → hand-written disclosure logic
       → HRIS / case system
```

Trust boundary: **the prompt and your custom guardrail code.** Attack surface = every field your integration code can pull.

Magic Cloudlet + Claude

```
Question → Claude
    → execution tree
       → per-role capability check → HRIS / case system
```

Trust boundary: **the runtime.** Attack surface = only the fields this role's grant explicitly allows.

A Magic Cloudlet exposes your HRIS and case-management system to Claude as a set of generated, role-gated endpoints — with each data field carrying a classification (operational, medical, disciplinary, compensation) rather than sitting in one undifferentiated case record. The agent's token carries a role, the same way a badge does. A manager role sees dates and coverage needs; an HR generalist role sees more; medical detail is scoped to HR and the employee, regardless of who's asking or how reasonably.

03 WHAT IT TAKES TO GET THIS TO PRODUCTION

Thirteen steps between an HR-copilot demo and one you'd trust with real cases

This is the checklist a security-conscious HR leader should put in front of a CFO before letting an AI agent answer questions against live employee records.

#	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 — zero migration.
2	**New-hire provisioning workflow**	MANUAL Built and maintained by your engineering team.	HANDLED Generated as endpoints — accounts, equipment, paperwork routing.
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 from 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 prompt hardening — a defense each new question tests.	HANDLED A reasonable-sounding question still can't exceed the capability 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 still reviewed.
11	**Compliance scope (ADA, FMLA, GINA, EEOC, state leave law)**	MANUAL Ongoing, and specific to your jurisdictions.	PARTIAL Self-host/air-gap keeps infra in your boundary; scope mapping is still yours.
12	**Hosting, deploy, rollback**	MANUAL A DevOps project, built once a backend exists.	HANDLED Included with a managed cloudlet, or self-hosted / on-prem / air-gapped.
13	**Ongoing patching & maintenance**	MANUAL Your team, indefinitely, as each HRIS API changes.	PARTIAL Maintained by Nuity (managed), or self-maintained under the MIT license.

```
13 / 13 MANUAL     Claude alone — every disclosure rule is code your team writes, tests, and re-tests as
                   roles and cases evolve
9 / 13 HANDLED     Magic Cloudlet + Claude — generated and enforced by the runtime
4 / 13 PARTIAL     review scope shrinks to case-routing logic and your specific jurisdictional compliance
                   policy
```

04 SECURITY ARCHITECTURE, LAYER BY LAYER

Five places an over-disclosure can be stopped — or not

A security reviewer's real question isn't "is the model discreet." It's "what, mechanically, refuses to disclose the field this role isn't entitled to see" — no matter how legitimate the request sounds.

#	LAYER	CLAUDE ALONE, FREE-FORM	MAGIC CLOUDLET + CLAUDE
1	**Field vocabulary**	Open — any field your integration code can read, reachable if the prompt gets there.	Closed — a finite, classified set of fields per role.
2	**Pre-execution verification**	Prompt hardening and application-level checks, written by hand.	A static proof against the live capability registry, before delivery.
3	**Capability grants (data classification)**	Not native — hand-built into the case-summary functions.	Argument-bound grants, checked by the evaluator at every dispatch.
4	**Identity & role (manager vs HR vs legal)**	A middleware pattern each new build may implement differently.	A runtime object — one role, from the caller's token to the admin screen.
5	**Audit trail**	Whatever your team designs, instruments, and maintains.	Built-in, severity-tagged log of every access and disclosure.

THE LINE A SECURITY TEAM CARES ABOUT

Everything on the right is enforced at the moment of dispatch, by the evaluator, on every call — not by a prompt a well-meaning manager's reasonable question is free to talk past. That difference is architectural, not behavioral. It doesn't degrade as case volume, or good intentions, go up.

The platform bill is the small number

	Claude alone + your team	Magic Cloudlet + Claude
Platform cost	No separate platform fee — cost shows up in engineering time and ongoing disclosure-logic maintenance.	Self-hosted: free, MIT-licensed. Managed: flat **$100/mo**, one developer user included.
Build cost, reference integration	~140,000 tokens of generation, debugging, and hardening the disclosure/routing logic.	~25,000 tokens — roughly an **80% reduction** on the backend/integration slice.
Time to a governed copilot	Days to weeks of review and hardening, after the demo has already shipped.	Minutes to scoped, role-gated endpoints; a working session to the full agent.
Ongoing maintenance	Your team, indefinitely, as each HRIS API and each new case type evolves.	Included on managed plans, or self-run under the MIT license.

What getting this wrong actually costs

These figures are the EEOC's, not Nuity's — from the agency's own fiscal-year 2025 enforcement data.

$660M
recovered by the EEOC for 17,680 workers in FY2025 — the agency's own enforcement total for the year.

$40K
average out-of-court discrimination settlement, per the EEOC — before a case ever reaches trial.

2nd
most common basis for EEOC lawsuits in FY2025 was disability discrimination — exactly the category a mishandled leave case falls into.

Source: U.S. Equal Employment Opportunity Commission, Fiscal Year 2025 enforcement data. Figures are agency-wide totals and averages, not a Nuity AI estimate or a claim about any specific customer's risk.

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.

06 WHERE WE'RE HONEST ABOUT TRADE-OFFS

What Claude alone still does better

A brief a budget approver can trust has to concede real points, not just make them.

GENUINELY TRUE, TODAY

Claude alone is faster to a rough internal prototype — a drafting assistant that helps an HR generalist write an offer letter or a policy summary, with no read access to case records at all. If your first step is "let's see if the model even understands our onboarding paperwork" and nothing it does can touch a real employee's file, hand-rolling that pilot is a completely reasonable way to spend an afternoon.

What no amount of prompting changes

- **The question is still legitimate — that's the problem.** Prompt-level filters are tuned to catch bad intent; a manager's genuine, well-meant question doesn't look like an attack, so it doesn't get caught.
- **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.
- **Review burden doesn't shrink because the model is good.** Free-form disclosure logic still needs a human compliance pass before real employee data hits it.

THE ONE-SENTENCE VERSION FOR YOUR APPROVER

Claude is the best available way to understand what someone's really asking. A Magic Cloudlet is what makes sure the answer only includes what that person is legally entitled to know — no matter how reasonable the question sounded.

07 NEXT STEP

Twenty minutes. Bring your real HRIS.

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

SELF-HOSTED

Free

MIT-licensed runtime — audit it, fork it, run it behind your own firewall.

MANAGED CLOUDLET

From $100/mo

Hosted, secured, and kept current — 1 developer user included.

MANAGED SERVICES

From $3,000/mo

A dedicated Magic/Hyperlambda developer on your team, 10 hrs/mo.

Nuity AI

Dallas–Fort Worth, TX

clayton.redmon@nuity.ai · nuity.ai

© 2026 Nuity AI

Magic Cloudlets run on Hyperlambda — MIT licensed

— METHODOLOGY & SOURCES

The fine print, up front

In the same spirit as the rest of this brief: here's exactly where the numbers above come from, and where their edges are.

- **Token and cost figures** are Nuity AI's published estimates as of mid-2026, based on a reference build (two linked tables, eight CRUD endpoints, access locked to a single role). They reflect the backend/integration slice of engineering work only. Full methodology: `nuity.ai/savings-calculator`.
- **EEOC figures** ($660M recovered for 17,680 workers, ~$40K average out-of-court settlement, disability as the second most common FY2025 basis) are from the U.S. Equal Employment Opportunity Commission's own fiscal-year 2025 enforcement data and public reporting on it. These are agency-wide totals, not a prediction for any specific company, and not a Nuity AI estimate.
- **The manager-question example** on page 3 is illustrative, built to demonstrate the mechanism — it is not a real employee case or a documented incident.
- **This use case** — HR onboarding and case management — is a Nuity AI application of the Magic Cloudlet architecture to a common enterprise AI initiative. It is not one of the build shapes currently published on nuity.ai/use-cases; treat it as a proposed pattern, not a shipped case study, until a specific build exists to point to.
- **This brief is not legal advice.** ADA, FMLA, GINA, and state leave-law obligations vary by employer size, jurisdiction, and fact pattern; consult employment counsel before relying on any specific disclosure rule.
- If a figure here looks wrong to your team, we'd rather hear about it before your finance partner does — `clayton.redmon@nuity.ai`.