# Law, Policy, Culture, Purpose

How HR agents move from compliant to right

Kerri O’Neill and Venkata Pingali

THE CORE IDEA

# The goal is not a fixed HR agent, but a learning system that becomes increasingly aligned with the organisation it serves.

Law and policy tell an agent what it must and may do. Culture and purpose tell it how and why. Getting all four right is continuous: make them explicit, resolve the tensions, then observe, govern and teach.

THE PROBLEM

# Compliant is not the same as right

An employee whose mother is terminally ill asks to reduce his hours for six months to care for her. An HR agent handles the first response.

A compliant reply is easy. The right one depends on the law, what the organisation offers, how it treats people, and what it exists to do.

### Agents are moving from answering to acting

Changing records, starting workflows, recommending consequential decisions.

### Law-only agents default to the floor

Treating compliance as one block of rules misses three of four layers.

### Explicit, but not rigid

Codified knowledge must still tolerate ambiguity and know when to involve a person.

THE FRAMEWORK

# Four layers an HR agent must work across

MUST

## Law

The non-negotiable floor: what the agent must and must not do.

MAY

## Policy

What the organisation chooses to do, within and above the law.

HOW

## Culture

How things really work here: norms, tone and unwritten rules.

WHY

## Purpose

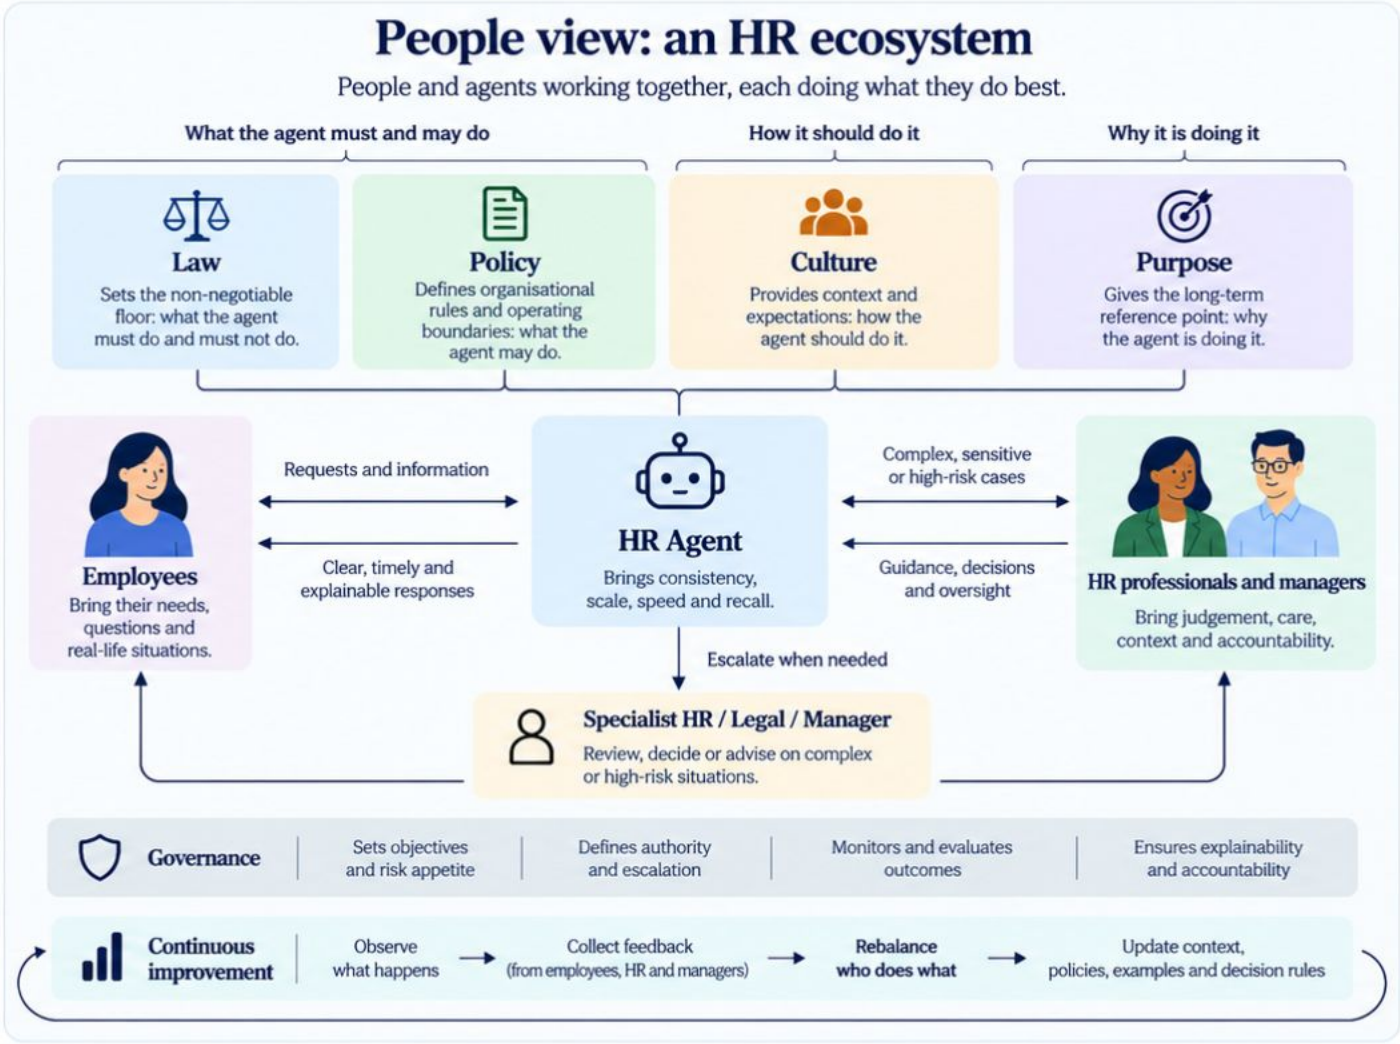
Why the organisation exists: the reference point when rules run out.

Onboard agents like new employees: as learning systems, not fixed software. Which raises a question HR cannot leave to others: **who teaches the HR agents?**

PEOPLE VIEW

# People and agents, each doing what they do best

The agent brings consistency, scale and recall. People bring judgement, care and accountability, and take the complex, sensitive or high-risk cases.



LAYER 1 · LAW

# The non-negotiable floor

16

distinct classes of legal instrument surfaced by mapping regulatory requirements

Risk surface = **jurisdictions × instruments × HR activities**

The task is not to give an agent access to the law, but to tell it what matters for **this person, this action, this jurisdiction, at this point in time.**

Preserve ambiguity, provenance and dependencies. Don’t manufacture precision where the law requires judgement. Where requirements conflict or stakes are high, involve a person.

LAYER 2 · POLICY

# What the organisation chooses

### Above the floor

Enhanced leave, hybrid working, a disciplinary process that starts with a conversation.

### Skipping it is a legal risk

UK tribunals look closely at whether employers followed their own policies.

### Precise and versioned

Which policy, which version and which provisions apply to this person, today.

No policy covers every case. The agent must understand what a policy is **for**, not just what it says.

LAYER 3 · CULTURE

# How things really work here

## “AI amplifies culture; it does not correct it.”

### Compliant can still be wrong

A formal warning where the culture expects a conversation first breeds mistrust.

### Avoid the average

Out of the box, agents sound like every other employer. That hollows out connection.

### Teach the intended culture

Principles, good and bad examples, and a clear list of what you would never do.

LAYER 4 · PURPOSE

# Why the organisation exists

Law, policy and culture look back. Purpose points forward.

### When rules run out

Guides action when no policy fits or the situation is new.

### Sharper risk judgement

Tells bending a rule apart from breaking its intent.

### More than a statement

Taught through choices and hard cases; never overrides law.

SCENARIO

# One request, four layers

## Law

UK: right to request flexible working from day one. US: FMLA may allow 12 weeks of reduced hours, short of six months. His mother's diagnosis is sensitive data.

## Policy

The employer may offer far more: paid carer's leave, or a temporary arrangement with a review date. He may not know these exist.

## Culture

Here, his manager may call him before any form is sent. The right move may be handing over to a person.

## Purpose

If the organisation exists to help people do their best work over a career, this is an investment, not a concession.

TRUST

# The person on the other side

He won’t experience a layered architecture. He will experience a reply, at one of the hardest moments of his life.

Does he know he is dealing with an agent?

Does it sound like his employer at its best?

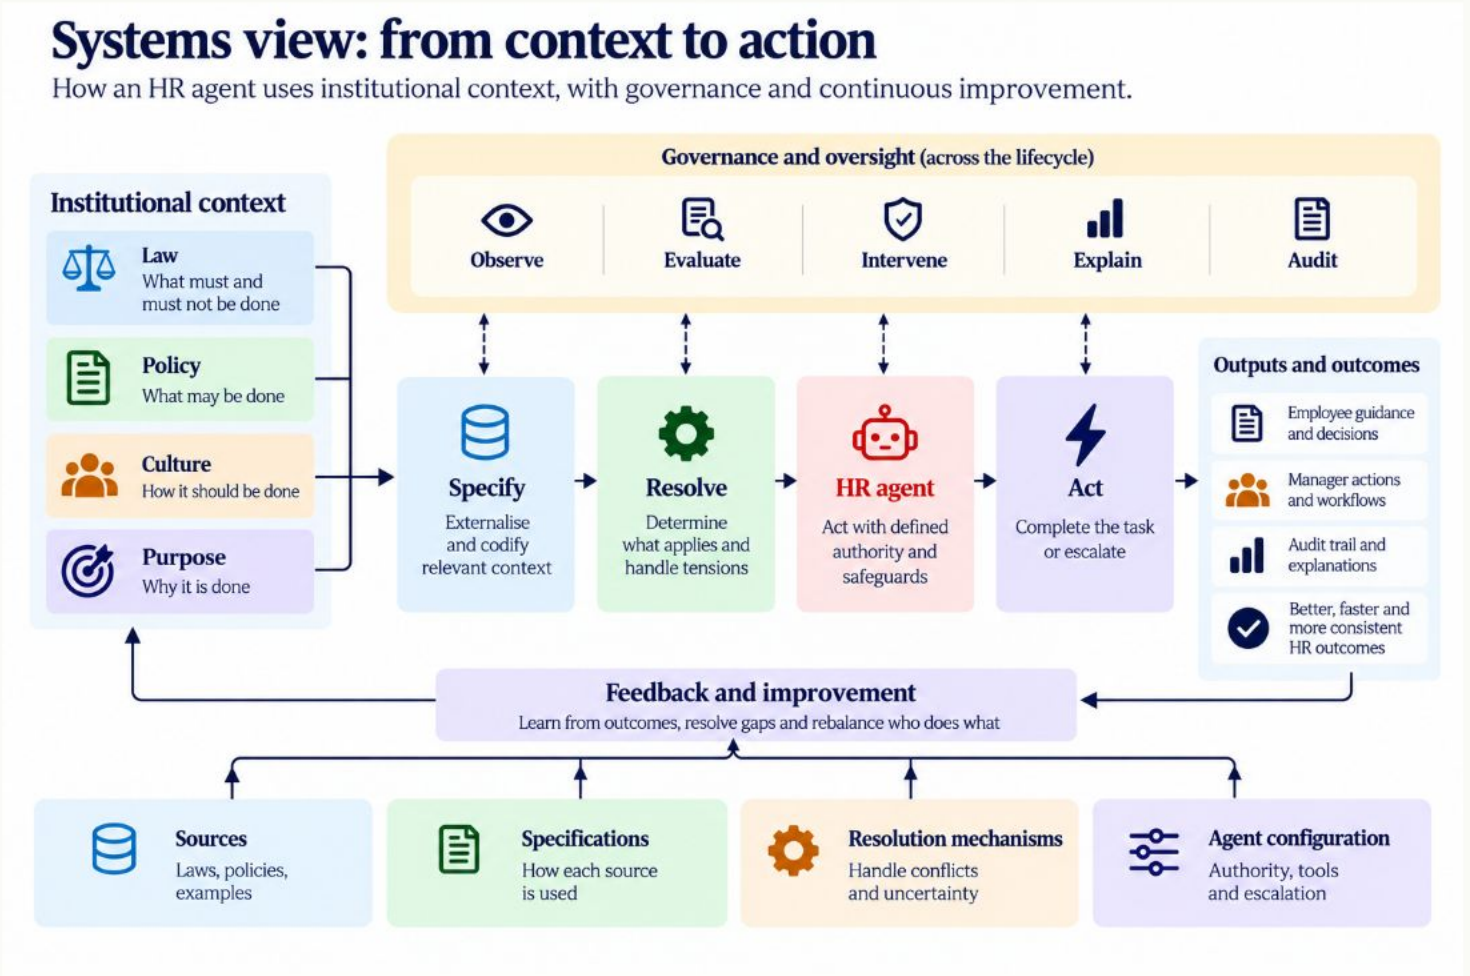
Can he reach a person easily?

Transparency is a design principle, not a disclaimer. Agents don’t change how trust is built; they multiply the interactions in which it is won or lost.

SYSTEMS VIEW

# A system around the agent

The layers can't just go into a prompt. Law becomes checkable obligations by jurisdiction and date; policy is versioned against law; culture and purpose shape judgement and escalation.



GOVERNANCE

# Governed, and built to learn

As ambiguity or consequence rises, autonomy falls and human involvement rises.

### Governance surrounds the process

Before acting: constrain, require approval or prevent. After: observe, evaluate, explain and audit.

### Limits are why it must learn

Mapping errors, tacit culture and changing law are certain. Every outcome should improve the sources, specifications and agent.

WHAT HR SHOULD DO NOW

# HR teaches the agents

Legal owns the law layer. HR owns policy intent, intended culture and purpose. The Chief People Officer is accountable for **how** HR agents behave, not just whether HR uses them.

01

### Write your “never” list

What would you never do to an employee, beyond what the law forbids?

02

### Build a library of good and bad

Real letters and decisions at your best and worst.

03

### Define authority before autonomy

What an agent does alone, recommends, or hands to a person.

04

### Keep accountability human

Responsibility for consequential decisions cannot be delegated.

THE QUESTION FOR HR

# Are we prepared to make explicit what we have always left implicit, and to own how our agents learn to act on our behalf?

We welcome examples: what have you made explicit, what resists codification, and where does human judgement still matter most?

Kerri O'Neill · Venkata Pingali