Policy enforcement at the rendering boundary

Govern what is shown and what is allowed back into the application

Most security controls do their work before the session starts: identity checks, device posture, conditional access, and network controls. Those matter, but they do not answer the question that increasingly determines whether agentic automation is governable in production: what is allowed to happen once the session is underway?

This checklist is designed to help security teams evaluate whether an approach can enforce policy at the rendering and interaction boundary, where interfaces are displayed and where user or agent actions flow back into systems of record.

Use this checklist when:

- Reviewing a new agent workflow before production approval
- Comparing platforms or architectures for governed agent execution
- Validating whether a pilot can be safely expanded

Checklist (policy enforcement)

- ☐ **Pre-render control exists**: Can you prevent unauthorized content from rendering, not only watermark or detect after the fact?
- ☐ **Granularity is meaningful:** Can policy scope be narrower than an app, for example a field, element, record, section, or document region?
- ☐ **Decision inputs are explicit:** Are decisions based on identity, role, workflow state, and data class, plus session risk signals where relevant?
- ☐ **Bidirectional policy:** Are both "show" and "do" governed, including copy, paste, download, print, upload, and high-risk UI actions?
- ☐ **Deterministic enforcement:** Is enforcement consistent across endpoints, without relying on local agents or extensions to behave correctly?
- ☐ **Policy change control:** Is there an audit trail for policy edits (who changed what, when, why), with rollback?
- ☐ **Break-glass behavior:** Is there a controlled, auditable exception path for emergencies, with strict scope and time limits?
- ☐ **Session boundaries are clear:** Can you scope policy per workflow or per session, rather than creating permanent broad access?
- ☐ **Evidence is produced:** For each enforcement decision, can you log the policy evaluated, inputs, and outcome for later review?