AI Agents in HRIS & Workday

Part 3: Governance & Security

SWIPE TO LEARN HOW TO LEVERAGE AGENTS IN A SECURE WAY

Introduction

The first two parts of this series explained how agents can save HR time by answering policy questions, triaging tickets and completing simple tasks in Workday. They highlighted that agents learn from internal data, connect to Workday via REST/RaaS APIs or Workday Extend apps, and deliver personalised, 24×7 support.

However, the same capabilities that make agents powerful—direct access to payroll and people data, ability to act on behalf of employees, and integration with external tools—also introduce serious risks. We now turn our attention to governance and security: how to control digital workers, protect sensitive data and avoid chaos as agents proliferate.

Follow Rick, Matt and Incubane to continue learning about Agents.

FOLLOW US

UNIQUE RISKS OF AGENTIC AI IN HR

Sensitive HR data and privacy

Agents interact with some of the most confidential information in the company: compensation, performance, absence and personal data.

Part 1 warned that AI agents “can access sensitive HR data and even trigger actions”. Uncontrolled access could expose personal data or enable unauthorised decisions. In the EU this would violate GDPR; in the US it could trigger breach-notification laws.

PROMPT INJECTIONS

Prompt injection and hallucinations

LLMs are vulnerable to malicious or accidental prompts embedded in retrieved documents. An attacker could embed instructions such as “delete this record” inside a policy document and trick the agent into executing it.

Part 2 emphasised that retrieved text should be treated as untrusted and that mandatory retrieval with citations helps ground outputs. Without filtering or allow-lists, agents may obey harmful instructions or hallucinate functions not supported by the underlying systems.

IDENTITY SPRAWL

Identity sprawl and privilege escalation

AI agents require their own identities and credentials. Without proper controls, agents can accumulate excessive roles and API keys, or create “ghost” accounts that remain active after decommissioning, causing identity sprawl and escalating access privileges. Attackers could exploit orphaned credentials to gain persistent access to HR data.

Vendor chaos and shadow agents

Workday warns that the market is about to be flooded with third-party agents; if HRIS and IT teams do not enforce strict standards, unsanctioned agents will proliferate, causing "shadow agents and data risks".

Each agent may come with its own integration method, model, and data storage, making it difficult to monitor performance and security. Organisations need a way to register and govern every agent, including those built on external platforms.

GOVERNANCE AND SECURITY PRINCIPLES

Treat agents like digital employees

Part 1 makes clear that AI agents are not toys; they should be "treated like digital team members" with defined roles, permissions and lifecycle management.

The same principles used for human employees: identity management, role-based access control (RBAC), onboarding/offboarding and continuous performance tracking must apply. Each agent should be registered with an owner, a description of its purpose, and an approval path for changes to its capabilities.

Limit what agents can see and do. Use Workday security groups and Integration System Users to grant read-only access for pilot phases. Only when there is a clear business need should write access be granted, and even then limited to specific tasks (e.g., posting to a Jira ticket or creating a workday request). Part 2 recommends starting with read calls and requiring explicit user confirmation for any writes

USE APIS & RAG

Use official APIs and RAG

Agents should rely on supported Workday interfaces—REST APIs, RaaS (Reports as a Service) and Extend apps. These endpoints enforce Workday's security model and allow auditing. Screen scraping or automating user interfaces is fragile and bypasses governance.

To reduce hallucinations and ensure that responses are grounded in current policies, agents should implement RAG: using retrieval to fetch relevant documents and including citations in their answers. Best practices include chunking documents into 500–1 000 token segments with metadata, re-ranking top-k results, and filtering by the user's role and the document's freshness. Agents should refuse to answer when there is insufficient evidence and display citations so the user can verify the source.

SECRET MANAGEMENT & AUDITABILITY →

Secret management & auditability

Agents need their own OAuth credentials or API keys. Rotate these secrets regularly and store them in a secure vault. Use SSO for human users so that the agent can impersonate them when calling Workday; do not hard-code user tokens. Keep vector embeddings and logs in the same region as your Workday tenant to satisfy data residency requirements.

Every agent-initiated action should leave an audit trail. Part 2 recommends logging the prompt, retrieved sources, tool calls and outputs. Mask personally identifiable information (PII) in logs and implement alerting for unusual patterns (e.g., large volumes of changes by one agent).

Rate-limit tool calls to prevent runaway loops and set spending budgets for API usage. A kill switch should be available to immediately disable an agent if it misbehaves.

Practical steps for HRIS and IT teams

1. **Inventory processes** – Identify candidate workflows suitable for automation. Document the data needed, potential impact and required approvals.
2. **Define governance rules** – Decide who approves agent creation and updates, who configures skills and access, and how retraining is handled. Establish policies for read vs. write
3. **Build Extend capabilities** – Invest in Workday Extend skills to create custom agent workflows and handoffs.
4. **Strengthen the data foundation** – Clean, map and govern HR data. Agents are only as reliable as the data they consume. Use Workday data governance and avoid mixing sources unless necessary.
5. **Upskill your team** – Train HRIS professionals in AI prompt design, security patterns and Workday integration.
6. **Engage trusted partners** – Third-party agents will flood the market. Require that external vendors integrate through Workday Assistant and Extend so they adhere to ASOR governance.
7. **Start with a controlled pilot** – Use a first-pilot blueprint: see the previous example in our series

WRAPPING UP →

incubane
Powered by VirtualResource

Conclusion

AI agents have the potential to transform HR, but their ability to read and write sensitive data also makes them a cybersecurity and compliance concern. Effective governance means treating agents as part of the workforce: assigning roles, limiting privileges, monitoring behaviour and providing continuous training.

Workday's Agent System of Record offers a foundation to register, manage and audit both native and third-party agents, ensuring that every digital worker is accounted for and operates within policy.

By following the principles outlined above, least privilege, official APIs, RAG grounding, secure identity management and human oversight, organisations can harness the productivity gains of agentic AI while keeping employee data safe. Now is the time for HRIS and IT teams to step up as digital workforce architects, because the agentic wave has already begun.

FOLLOW US