



The AI Platform Checklist

A Practical Evaluation Framework for Enterprise IT Leaders



Enterprise
moving from
pilots to s
infrastructure
CIOs, IT leaders
AI councils, the
question is no l
Which AI tool



How to use this checklist

Enterprise AI is moving from isolated pilots to strategic infrastructure. For CIOs, IT leaders, and AI councils, the question is no longer “Which AI tool performs best?” It’s “Are we building a scalable enterprise AI platform, or accumulating long-term fragmentation?”

This checklist provides a structured framework to evaluate whether your AI approach:



Can scale beyond isolated chatbots and assistants



Integrates with your existing architecture



Provides centralized visibility and control



Aligns with your enterprise governance standards



Converts AI investment into measurable business impact

This checklist is designed to support:



AI council discussions



Vendor evaluations




Enterprise-wide standardization decisions



Executive and security stakeholder alignment





1. Core platform capabilities

 **Main question:**
Does this platform have the infrastructure required to scale enterprise AI value across your most critical workflows?


- Ubiquity**
- Is AI embedded across existing work surfaces (email, docs, CRM, ticketing, etc.), rather than confined to a stand-alone interface?
 - Can users access AI without switching tabs, copying context, or re-explaining their work?
 - Does the platform eliminate the need for a stand-alone AI destination?
 - Does it improve, refine, and extend the content users are already creating, rather than requiring them to paste work into a separate tool?

-
- Proactivity**
- Does the platform surface recommendations, drafts, or actions at relevant moments within workflows?
 - Can AI initiate assistance based on context rather than requiring explicit prompts?


-
- Connected intelligence**
- Can the platform securely connect to core enterprise systems across your stack?
 - Does AI operate with awareness of user, team, and organizational context?
 - Can it orchestrate workflows and take action across systems, not just generate outputs?

-
- Collaboration**
- Can AI operate across team workflows rather than being confined to individual sessions?
 - Does the platform allow teams to build on shared context and collective work?
 - Can insights, outputs, or agent behaviors be standardized and improved across departments?

-
- Advanced capabilities**
- Does the platform support specialized AI capabilities, like data analysis, coding assistance, or image and media generation?
 - Can teams build, deploy, and govern custom agents with appropriate administrative controls?
 - Can the platform integrate or extend via APIs to support evolving AI capabilities?
 - Is there a clear path to expand functionality without introducing new AI silos?

 **Red flag:**
If AI depends on users leaving their workflow, writing detailed prompts, manually stitching together tools, or operating in isolated sessions, adoption will stall and platform value will plateau.





2. Platform architecture

 **Main question:**
Are we implementing a scalable AI platform or accumulating disconnected tools?

Architecture and orchestration


- Is there a centralized orchestration layer across AI tools and models?
- Can we support multiple LLMs without re-architecting workflows?
- Does the solution avoid vendor lock-in to a single model or ecosystem?
- Does the platform provide mechanisms (e.g., APIs, SDKs, or admin controls) to deploy and manage agents at scale?

Enterprise integration

- Does it integrate with our existing identity provider (SSO, SCIM)?
- Does it inherit existing role-based permissions and access controls?
- Can it integrate cleanly with enterprise systems without heavy customization?
- Is deployment achievable without extended IT lift?


Scale and standardization

- Can this support enterprise-wide deployment?
- Does it reduce tool sprawl, rather than introduce another AI surface?
- Is there centralized admin visibility across business units?

 **Red flag:**
If every department is piloting disconnected AI tools, you don't have a platform; you have fragmentation.



3. Governance



Main question:
Can governance be enforced at the platform level?

Administrative control

- Are there centralized admin controls for AI capabilities?
- Can capabilities be configured by role, team, or department?
- Are AI usage logs accessible for audit and compliance review?
- Is there visibility into usage across business units?


Policy enforcement

- Can enterprise policies be embedded directly into the platform?
- Can we control which data sources AI can access?
- Are guardrails configurable for sensitive workflows?
- Can disclaimers, compliance language, or brand standards be enforced?



Red flag:
If governance depends on training alone rather than enforceable controls, risk will scale with adoption.





MS or bioavers
pporting thou
pecies of fish, m
d crustaceans.
osystems not c

4. Security, privacy, and risk controls



Main question:

Are security and risk controls embedded in the platform architecture, or layered on after deployment?

Data handling and privacy

- Is enterprise data excluded from model training by default?
- Are data retention policies transparent and configurable?
- Is data encrypted in transit and at rest?
- Are regional data residency requirements supported?

Risk mitigation

- Are there protections against prompt injection and data leakage?
- Is sensitive data detection and redaction available?
- Are role-based access controls enforced consistently?

Compliance

- Are enterprise certifications (SOC 2 Type II, ISO 27001) current and independently validated?
- Is the platform aligned with relevant privacy regulations (e.g., GDPR)?
- Are independent penetration tests conducted regularly?
- Is vendor risk documentation readily available for review?



Red flag:

If security documentation is incomplete or unclear, assume risk will surface later in procurement.





Interpreting your results

Your responses should reveal whether you are evaluating AI as a feature set or as long-term enterprise infrastructure. Consider the patterns in your answers:

If you consistently answered yes, you are evaluating AI as a platform, not a tool.

If you identified gaps in core capabilities, you may be evaluating a solution that improves productivity individually, but does not function as scalable infrastructure.

If governance or security responses were inconsistent, you will introduce risk and fragmentation that will compound as adoption scales.

Strategic reminder

Platform decisions compound. The AI solutions you standardize today will shape your governance model, vendor landscape, and ability to scale connected, collaborative intelligence across teams.

About Superhuman

Superhuman (formerly Grammarly) is the AI productivity platform on a mission to unlock the superhuman potential in everyone. The Superhuman suite of apps and agents brings AI wherever people work, integrating with over 1 million applications and websites. The company's products include Grammarly's writing assistance, Coda's collaborative workspaces, Mail's inbox management, and Go, the proactive AI assistant that understands context and delivers help automatically. Founded in 2009, Superhuman empowers over 40 million people, 50,000 organizations, and 3,000 educational institutions worldwide to eliminate busywork and focus on what matters. Learn more at superhuman.com/solutions/enterprise.