Industry AI Employees: Core Problems, First Tasks and Acceptance Evidence Across Eight Sectors

Design by sector risk, data, workflow and accountability instead of covering every industry with one chat assistant

Online technical whitepaper

6 min read

Published / updated
2026-11-22

Intended readers

Industry business leaders, digital transformation teams, security and compliance professionals, solution leaders and procurement teams

EXECUTIVE ANSWER / DIRECT ANSWER

An industry AI program should not begin with a generic assistant. It should identify a valuable, verifiable core problem. Legally authorized defense-related assistance and enterprise protection emphasize authority and evidence; government and finance emphasize accountability and definitions; manufacturing and retail emphasize system coordination; research and office work emphasize knowledge quality. Each requires its own tasks, permissions, human ownership and acceptance evidence.

01 / KEY QUESTIONS

Key questions answered in this paper

Q1

Why cannot all eight sectors use one acceptance template?

Error consequences, data sources, process points and regulatory duties differ. The platform can be shared, but evaluation sets, thresholds, permissions and sign-off owners must be sector-specific.

Q2

How should the first industry task be selected?

Choose frequent, bounded work with inspectable output, reversible actions and historical examples. Begin with analysis and drafting, then expand only after validation.

Q3

Can AI directly make high-risk interventions or financial decisions?

Not by default. High-risk interventions, device control, financial judgment and other consequential actions require authorized accountable people and project security and compliance assessment.

02 / Full paper

01

Core problems and first acceptable tasks in eight sectors

Every sector needs a problem–task–evidence loop before a feature list.

The table below is a solution-design starting point, not a claim about customer outcomes or completed cases. Actual tasks require project discovery of data, systems and responsibilities. Start with inspectable research, analysis and drafting before considering limited system actions.

China’ s AI Plus initiative emphasizes integration with sectors and scenarios, while CAICT industry research examines coordination across technology, industry and governance. Enterprises should translate these directions into their own task portfolios and capabilities; macro policy is not a guarantee of project returns.

Sector	Core problem	Example first task	Key acceptance evidence
Defense-related, legally authorized assistance	Distributed material and heavy knowledge and compliance work	Authorized document organization, policy knowledge support and exercise-record drafting	Authorization scope, source traceability, human review and audit logs
Enterprise protection	Fragmented alerts and knowledge with inconsistent handoffs	Defensive alert summaries, response-knowledge retrieval and report drafting	False-positive/miss review, least privilege and analysis–action separation
Digital government	Complex policy definitions and accountable responses	Policy retrieval, assisted material review and response drafts	Current authority, consistent definitions, human issuance and personal-data protection

Sector	Core problem	Example first task	Key acceptance evidence
Quantitative finance	Separated data and research with difficult reproducibility	Research collection, factor-code assistance and backtest-report drafting	Point-in-time data, code review, reproducibility and human investment responsibility
Smart manufacturing	Isolated equipment knowledge and quality/maintenance gaps	Maintenance knowledge support, quality-exception summaries and shift handover	Time-aligned data, engineering review and preserved equipment safety controls
Technology intelligence	Broad sources, uneven evidence and repeated tracking	Technology and patent research, competitive summaries and research drafts	Real sources, evidence coverage and separation of fact from inference
Commerce & retail	Fragmented product content and fast campaign/analysis cycles	Product-material organization, content drafts and business-analysis summaries	Product-fact checks, brand approval, data definitions and publication authority
General office	Cross-team information gaps and repetitive documents and meetings	Meeting notes, knowledge Q&A, reports and process drafts	Participant confirmation, knowledge permissions, quality and rework records

Chapter sources [1] [2] [3]

02

Reuse the foundation; build sector assets separately

Runtime and governance can be reused; knowledge, evaluation and accountability cannot simply be copied.

VROS Agent can provide common application, event, memory, space and Worker-sandbox mechanisms. AIBotStore organizes role capabilities, multiple compute forms host models and tools, and FDE maps them into the field. Model, runtime, sector knowledge, business systems and final responsibility must be delineated rather than attributing every capability to one model.

Sector assets include terminology and policy bases, task templates, tool connectors, evaluation sets, risk inventories and playbooks. Each needs an owner, version and scope. Events, permissions, audit and evaluation methods can transfer across sectors, but sector data, failures and thresholds must be rebuilt.

- ✓ Common layer: task state, identity, events, memory, sandboxes, monitoring and versions.
- ✓ Sector layer: knowledge, terminology, data, workflows, tools, evaluation and compliance boundaries.

- ✓ Organization layer: accountability, approval, training, takeover, change and supplier management.
- ✓ Infrastructure layer: model serving, compute, networking, storage, backup and continuity.

Chapter sources [3] [4] [5]

03

Global trends: from model capability to action governance

Adoption now faces interoperability, post-deployment monitoring, risk management and transparency together.

NIST’s AI Agent Standards Initiative includes agent identity, protocols and security research. NIST AI RMF and AI 800-4 address risk management and post-deployment monitoring. The trend is not that complete autonomy has become the default; as agents act, identity, permissions, evaluation and operating evidence become more important.

The EU AI Act uses a risk-based approach and includes transparency and high-risk-system requirements. Cross-border organizations need to assess obligations by product, role, user and jurisdiction. Referencing a framework is not universal compliance. China’s generative-AI and content-labeling rules likewise require assessment against actual service and publication scenarios.

- ✓ Technology: long-horizon work, tool use, multi-agent systems and interoperability.
- ✓ Governance: identity, least privilege, supply chains, post-deployment monitoring and incidents.
- ✓ Compliance: risk tiers, transparency, content labeling and records.
- ✓ Procurement: real-task evidence, change and exit rather than feature lists alone.

Chapter sources [8] [6] [9] [11] [13] [14]

04

High-impact settings separate assistance from accountability

Systems may produce recommendations and evidence; authorized people remain responsible for consequential judgment.

Legally authorized defense-related work is limited to knowledge, material and process assistance and does not involve weapons or offensive capabilities. Enterprise protection may support defensive analysis and reporting, while high-risk intervention needs authorized confirmation. Quantitative finance may support research and code, but generated content is not unattended investment advice or an unowned trading decision.

External government responses, manufacturing-device control, resource allocation and other actions affecting people or production safety also require human issuance or an engineering safety chain. AI output is not a substitute for legal, investment or safety advice; obtain appropriate professional review for the actual scenario.

Boundary	AI contribution	Human responsibility retained
Knowledge & research	Retrieval, summaries, comparisons, drafts and sources	Source applicability and final conclusions
Business judgment	Options, risks and data basis	Authorized decision and sign-off
External actions	Prepare or execute bounded low-risk actions	Approval of high-impact or irreversible actions
Exceptions & incidents	Alerts, pause suggestions, records and escalation	Emergency response, notification and accountability

Chapter sources [4] [5] [6] [12]

05

Prioritization and acceptance include the cost of error

The same accuracy has different meaning under different consequences; acceptance must cover impact and recovery.

Prioritize by business value, data readiness, output inspectability, action reversibility, integration complexity and regulatory risk. Begin with human-reviewable assistance and build failure examples and a human baseline. Deepen automation only after quality, permissions, exceptions and organizational responsibility pass validation.

Acceptance reports should identify sector, task, sample period, model and application versions, knowledge scope, conditions, human baseline and uncovered cases. Retain passes, failures, rework, refusals, intervention and recovery. Individual demos, theoretical targets and third-party benchmarks cannot replace real project results.

- ✓ Build routine, difficult, conflicting and adversarial evaluation sets per sector.
- ✓ Include impact, discoverability, reversibility and review cost in risk tiers.
- ✓ Accept retrieval, analysis, delivery and actions by stage, not just final text.
- ✓ Reassess model, policy, knowledge and tool changes with rollback retained.

Chapter sources [7] [9] [4] [10]

06

Executive roadmap: from sector problem to reusable capability

Treat industry AI as a business-change and governance program, not a point software purchase.

First, business leaders define the core problem and red lines. Second, a cross-functional team selects the first task and baseline. Third, a controlled PoC combines VROS Agent, AIBotStore, compute and FDE. Fourth, evidence determines expansion. Contracts need data, version, asset, support, change and exit terms.

Brand communication follows the same evidence discipline: distinguish public capabilities, project configuration, evolution and third-party material without inventing customers, cases, certifications or results. ISO/IEC 42001, NIST and other frameworks may organize management questions; referencing them does not imply AIBotStore certification.

- ✓ Business: select the problem, task owner and unacceptable risks.
- ✓ Technology: establish governed task chains and real-environment compatibility evidence.
- ✓ Delivery: produce sector assets, training, acceptance and operating handover.
- ✓ Brand: communicate expertise through verifiable methods and evidence, not inflated outcomes.

Chapter sources [1] [3] [10] [6]

---

03 / CHECKLIST

Pre-project checklist

Use these items to structure a PoC plan, procurement questions or acceptance checklist.

01	Define core problem, first task and acceptance evidence per sector
02	Separate common runtime from sector knowledge, evaluation and accountability assets
03	Assess risk, transparency, content and compliance by jurisdiction and scenario
04	Retain authorized accountability for high-risk action, device control and financial judgment
05	Accept with real samples, failures, human baselines and recovery tests
06	Distinguish owned capability, project configuration, evolution and third-party sources in brand content

04 / FAQ

Frequently asked questions, answered directly

01 **Which of the eight sectors should deploy first?**

There is no universal ranking. Select from the enterprise’ s task frequency, data readiness, output inspectability, reversibility, risk and human baseline.

02 **Can one evaluation set be reused across industry solutions?**

Methods and some general security cases can transfer. Knowledge, terminology, difficult cases, thresholds and sign-off accountability must be rebuilt for each sector and task.

03 **Do the solutions provide unowned financial or high-risk security decisions?**

No. Financial judgment, high-risk intervention, device control and other consequential actions require authorized accountable people; AI performs assessed assistance or bounded execution only.

04 **Are the industry tasks in this whitepaper customer cases?**

No. They are solution-design and PoC examples, not claims of delivered customers or achieved outcomes. Actual projects depend on discovery, contract and acceptance.

05 / REFERENCES

References and public sources

Chinese references follow GB/T 7714 style and English references follow APA style. Accessed November 22, 2026. External material informs industry and technical methods; it does not imply endorsement or certification of AIBotStore.

[1] State Council of the People’ s Republic of China. (2025). Opinions on deepening the “AI Plus” initiative (Guo Fa [2025] No. 11).
https://www.gov.cn/gongbao/2025/issue_12266/202509/content_7039598.html →

[2] China Academy of Information and Communications Technology. (2026). Artificial intelligence industry development research report [Report].
https://www.caict.ac.cn/kxyj/qwfb/bps/202602/P020260202487301304903.pdf →

[3] China Academy of Information and Communications Technology. (2026). AI Native research report 1.0: Enterprise transformation framework and practice [Report].
https://www.caict.ac.cn/kxyj/qwfb/ztbg/202610/P020261022558735347833.pdf →

[4] China Academy of Information and Communications Technology, & China University of Political Science and Law. (2026). Agent governance research report (2026) [Report].
https://www.caict.ac.cn/kxyj/qwfb/ztbg/202609/P020260916574260931247.pdf →

[5] China Academy of Information and Communications Technology. (2026). Artificial intelligence security governance blue book [Report].
https://www.caict.ac.cn/kxyj/qwfb/bps/202601/P020260109784447548497.pdf →

[6] National Institute of Standards and Technology. (2023). AI Risk Management Framework (AI RMF 1.0).
https://www.nist.gov/itl/ai-risk-management-framework →

[7] National Institute of Standards and Technology. (2024). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1).
https://doi.org/10.6028/NIST.AI.600-1 →

[8] National Institute of Standards and Technology. (2026). AI Agent Standards Initiative.
https://www.nist.gov/caisi/ai-agent-standards-initiative →

[9] Rao, A., Keller, A., Kalra, N., Steed, R., Kwegyir-Aggrey, K., Klyman, K., Staheli, D., & Bergman, A. (2026). Challenges to the monitoring of deployed AI systems (NIST AI 800-4).
https://doi.org/10.6028/NIST.AI.800-4 →

[10] International Organization for Standardization. (2023). ISO/IEC 42001:2023—Information technology—Artificial intelligence—Management system.
https://www.iso.org/standard/42001 →

[11] European Commission. (2026). AI Act: Risk-based approach, transparency and high-risk system requirements.
https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai →

[12] OWASP GenAI Security Project. (2025). Agentic AI—Threats and mitigations.
https://genai.owasp.org/resource/agentic-ai-threats-and-mitigations/ →

[13] Cyberspace Administration of China, National Development and Reform Commission, Ministry of Education, et al. (2023). Interim measures for the management of generative artificial intelligence services.
https://www.gov.cn/zhengce/zhengceku/202307/content_6891752.htm →

[14] Cyberspace Administration of China, Ministry of Industry and Information Technology, Ministry of Public Security, & National Radio and Television Administration. (2025). Measures for labeling AI-generated and synthetic content.
https://www.gov.cn/zhengce/zhengceku/202503/content_7014286.htm →

Scope statement

This whitepaper supports technology selection and project planning and does not replace legal, compliance, safety or investment advice. Actual VROS Agent, AIBotStore, hardware-integration and controlled-evolution scope depends on the delivered version, project contract and acceptance. Industry task examples do not represent delivered customers or achieved outcomes.