When Your AI Agent has got to be right

kama.ai: Trust. Empathy. Accuracy.

kama.ai

AI's Next Big Phase

Responsible AI in Complex Tasks:
Enterprises moving Beyond Answer to Actions without Sacrificing Trust

https://kama.ai

Executive Summary

AI is entering a critical new phase, one defined not by how it talks, but by what it does. This is the phase of complex tasks. GenAI brought us powerful language tools. It made machines fluent, fast, and impressive. But it also introduced risk. Hallucinations, misinformation, and off-brand answers are now common in many deployments.

A 2024 Oxford study found that GenAI hallucinated in 58% of test cases.[1] McKinsey's 2025 report showed nearly half of all companies had experienced real-world consequences from GenAI errors.[2] The lesson is clear: fluency without governance is not enough. This is especially important as we step into the world of actions taken, with consequences.

Enterprises are now shifting focus, from answering questions to executing complex tasks. These actions require structure, oversight, and verified content. That's why responsible Hybrid AI is emerging as the most viable next step.

These systems combine deterministic Knowledge Graphs with governed GenAI and Robotic Process Automation (RPA). Together, they let enterprises automate multi-step processes safely, while maintaining legal, ethical, and brand standards.

This eBook explores why this matters, how Hybrid AI works, and where it applies. It also offers practical guidance on deployment and adoption.

Bottom line: Enterprises don't need less AI. They need AI they can trust. It's aboutI that acts with accuracy, transparency, and human oversight. That's the foundation of AI's next big phase.

GenAI Alone isn't Safe Enough

9%

ONLY 9 % OF SURVEYED ORGANIZATIONS ARE FULLY READY TO MANAGE GENERATIVE AI RISKS EFFECTIVELY.[7]

DELOITTE, 2025

GenAI has opened new doors, but also exposed us to new risks. It generates natural language, but not verified facts. It draws on training data, yet doesn't always cite sources. That makes it fluent, convincing - and *fallible*. For high-stakes settings, this is a problem.

Common issues are well known: hallucinated facts, off-brand tone, unreviewed output, offensive content, and biased sources. These aren't rare. Gartner's 2025 research identified nine types of 'Artificial Cognitive Disorders.'[3] These AI errors are expected to cost companies over $150 billion USD per year by 2030.[3] Yet McKinsey already found that nearly half (47%) of organizations have faced real consequences from GenAI mistakes, already.[2]

It's not just about what GenAI can do. It's about what it should be allowed to do. "The more we think about digital transformation and the application of AI," notes Mark Cosyn, VP at kama.ai, "the more important trust becomes, ... because we're replacing human judgment with machine actions."

None of us want an AI agent to make a mistake - like selling our investments at a loss. Accuracy and guidelines matter. Today, GenAI lacks this governance.

Despite this, companies don't need less AI. They need smarter AI - with controls, checks, and audit trails. The answer isn't to stop. The answer is to evolve.

Responsible AI gives you automation, creativity, and speed. This without compromising ethical or brand standards. The answers include Human-in-the-Loop workflows.

That's the shift. We're moving from uncertain answers to verified, actionable AI. That gives us a new foundation - one built on trust.

From Chat to Action

53%

OF BUSINESSES HAVE ALREADY IMPLEMENTED RPA TO MANAGE COMPLEX TASKS AND MORE PLAN TO ADOPT IT SOON. [8]

FLOBOTICS, 2025

The first wave of GenAI answered questions. You provided a question, it searched the web, and provided an answer, or crafted a message. The next wave is now focused on taking action. We need systems to manage complex tasks. Be it checking calendars to book a meeting, and sending a suggested agenda, or other more in-depth challenges - these are multi-step, context-sensitive, and compliance-driven needs.

HBR's research suggests that as we step into the action phase, AI will augment, automate, or reinvent up to 40% of all work activities.[3] This means this next phase will involve the significant work of interconnecting and automating many more systems to provide results from complex background interactions.

Complex tasks go beyond simple replies. Think onboarding a new employee, reviewing a contract, or solving a support issue. We need logic, policy alignment, and system integration. As kama.ai's *Hybrid AI eBook* points out, "Hybrid AI Agents don't just answer, they take actions. By connecting to Robotic Process Automation and governed GenAI, they can complete tasks across HR, Legal, and Customer Service while respecting enterprise guardrails."[5]

Here is where we need a new generation of AI. We need trustworthy Hybrid AI platforms that bring use both deterministic and guided probabilistic tools. These blend Knowledge Graph AIs with RAG based responsible AI and integrate RPA. The result isn't just a response. It's an action.

"The world does not need faster AI. It needs AI that is smarter, safer, and trustworthy."

Brian Ritchie, CEO, kama.ai

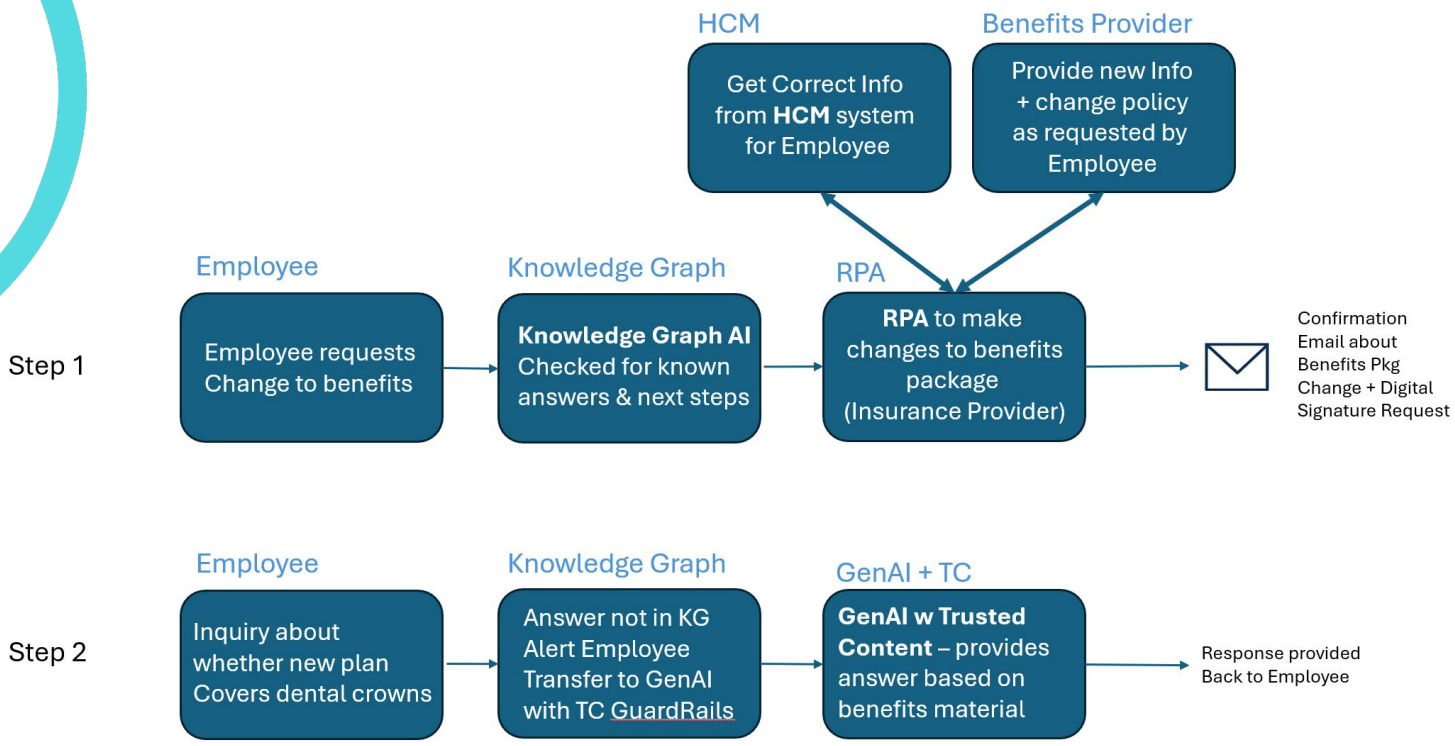

Imagine that a user needs to change their benefits package at work. They just got married, or had a baby, and need the family plan. The employee goes to the chatbot, to ask if it could make this change for them. With hybrid AI technology, the chatbot uses the knowledge graph AI to orchestrate the transaction. The knowledge graph knows this is a typical request. As shown in the diagram, it routes the need to the RPA (robotic process automation) system.

Since benefits updates are common, the RPA handles the transaction. The system authenticates the request and employee. It then retrieves important information about the individual from the HCM (human capital management) system. With this it can transact with the benefits provider to make the adjustment. Once done, it updates the HCM and payroll system as needed.

At the end of the sequence, the RPA also initiates an automated email to the employee. This updates them, about all the systems that have been adjusted, and requests their digital signature for final confirming approval.

But, the employee isn't quite done. Next they inquire with the chatbot about how much the coverage is for dental crowns under the new policy. Again the inquiry first goes to the knowledge graph. It turns out this information is not already in existence within the knowledge graph AI, so it is routed to the generative AI part of GenAI's Sober Second Mind®. Here the system uses a trusted collection (RAG) of all the benefits manuals and documentation. Now the generative side returns an answer but also alerts the employee that this answer was AI generated. Although it is highly likely to be correct and accurate - a confirmation with the documents is advisable.

An important point in these sequences is the use of different AI and automation technologies to accomplish complex tasks. In a global survey, 78% of organizations implemented or were planning to implement robotic process automation (RPA) solutions.[3] The importance here is to leverage, integrate and orchestrate new AI capabilities with existing, tested technologies that already solve certain problems.

Anatomy of Responsible AI

86%

86 % OF ORGANIZATIONS REPORT HIGHER PRODUCTIVITY AFTER ADOPTING RPA.[8]

FLOBOTICS, 2025

It starts with human-reviewed content. Subject matter experts and Knowledge Managers load verified answers into a structured Knowledge Graph. This deterministic layer delivers 100% accurate responses. As Brian mentions, “ask a question, if we have it in the knowledge graph, you’re going to get the absolute answer. You don’t have to wonder whether it’s right. You don’t have to check any underlying document links.”

This level of trust is key for customer-facing and high-risk use cases.

Next, documents from internal repositories are added to Trusted Collections. These are vectorized - making them searchable by meaning, not just keywords.

When a query arrives, the system checks the Graph. If there’s a match, it sends a verified reply. If not, GenAI drafts a response using only the Trusted Collections.

This draft:

- Goes to a Knowledge Manager for review, or
- Is sent to the user with a disclaimer

All actions are logged. Feedback improves the system over time - with human guidance. But, what makes this different from ChatGPT plugins or RAG tools?

First, containment: no open-web content is used.
Second, governance: GenAI is activated only when safe.
Third, accountability: all answers are traceable.

This separation of mission-critical and exploratory work matters. With deterministic safety, GenAI support, and RPA execution - all under governance - Trustworthy Hybrid AI lets you automate safely.

Enterprise Use Cases

27%

AI HALLUCINATION RATE IS 3%–27% IN ENTERPRISE DEPLOYMENTS. EVEN WITH PROMPT ENGINEERING, ENTERPRISE LLMS EXHIBIT ERROR RATES AS HIGH AS 27%. [2]

MCKINSEY & CO, 2025

Let's consider several use cases of how to use a Hybrid AI solution in a complex tasks scenario.

HR: Smarter Onboarding and Policy Access

An enterprise deploys Hybrid AI for onboarding. It guides new hires through policies, benefits, and tools. If asked about DEI or parental leave, the system provides verified answers. In other words, these answers are similar to what the new employee would receive, had they spoken with an HR representative. For particularly complex queries, it flags issues for HR to address directly.

In this case, the RPA sends reminders, logs queries, and books meetings. The benefit here is that HR is called on to handle fewer repetitive, mundane tasks. New hires get fast, accurate help, without stress or wasted time. Best of all, the new employee has a tool that helps them answer any and all personnel related questions, quickly, accurately, and without leaving them with a sense of being judged for silly questions.

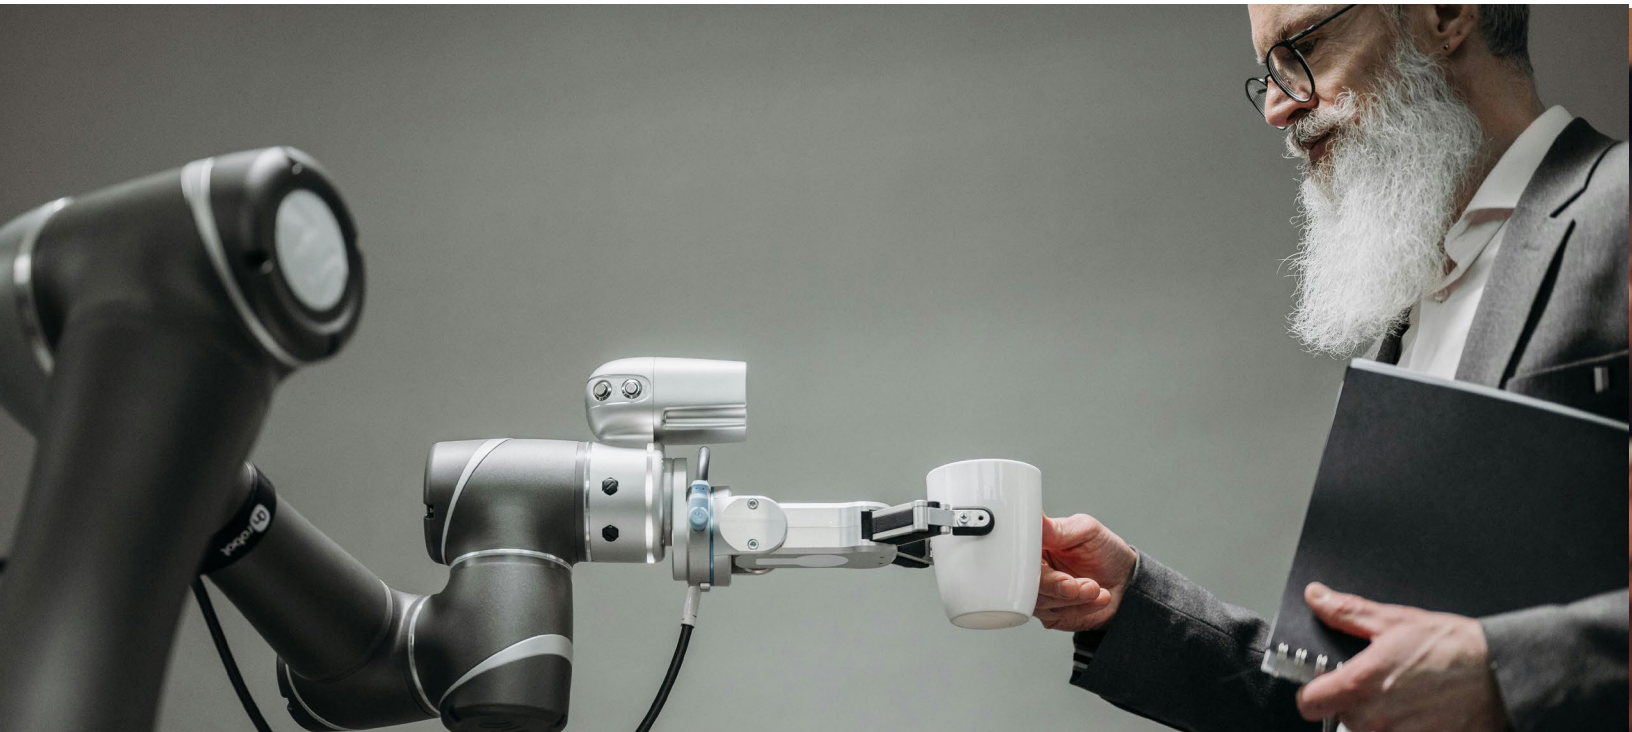

Legal: NDA and Contract Review

Legal teams use Hybrid AI to review NDAs. The platform scans documents, flags clauses, and highlights risk. GenAI can suggest edits which are promptly labeled as drafts which the legal SMEs review to either sanction as correct, or to adjust as the final decision. Given the sensitive and highly important nature of contracts, none of this information can be shared outside the organization. Merely using ChatGPT is not a good option as there is a chance of this information leaking to external sources.

By using a trusted AI solution like the Hybrid AI one suggested earlier, review times can drops by 40–60%. Standard clauses can be fully sanctioned and known by the system - for identification in contracts. Risk is lowered. Every edit is recorded. No outside content is used.

In this case the contract review AI acts not only as a reviewing mechanism, but can also be the contract repository of record for the organization. All contracts then being accessible for review by management and the legal team.

Customer Service: Smarter Support Flows

A telecom firm deploys the AI system on its public website. A customer asks about a billing error. Hybrid AI access the main billing system, and after using another RPA system for verification of the user's authenticity - provides details on the account. If the user is still not satisfied, the system offers step-by-step help or suggestions on options.

In this case, let's presume the user needs to fix an error with the user's information or billing details. If needed,

RPA creates a ticket and escalates the matter.

Since this is a customer engagement situation, brand tone, and brand voice need to be consistent. Verified answers coming from the sanctioned deterministic side of the AI (the knowledge base) are clear. If the user receives generative AI responses, they are labeled as such. In other words, the user is informed that these are not fully sanctioned answers, but there is a small chance of error with these, and that the user should use them with caution.

If the user proceed with these questions, finds an error, then the system can either use an RPA to connect with the billing / customer system to correct the error. Otherwise, a human agent handoff can engaged for more complex issues.

For each of these sample cases - the new option of complex task resolution is possible within a safe AI environment. Using the Hybrid AI - responsible AI technology lets an organization handle many such mundane tasks safely, with automation.

Deploying Hybrid AI

47%

OF EMPLOYEES BELIEVE AI WILL REPLACE 30% OF THEIR TASKS WITHIN THE FIRST YEAR OF DEPLOYMENT [2]

MCKINSEY & CO, 2025

From pilot to production should typically takes 2–3 months.

First, consult with AI experts and stakeholders. Set your architecture, KPIs, and adoption plans. Taking the step of speaking with experts at this stage and having an early consultation can save your organization considerable angst. Gartner estimates that, "over 80% of AI projects fail - twice the rate of other IT projects."[6] With this in mind, starting off with a consultation, and AI workshop is well worth the small costs to ensure your organization's AI project will succeed, and be adopted by your user base.

A good solution will start with content curation to create trusted content repositories. These are your company sanctioned documents. It becomes the basis for the Knowledge Graph AI.

Next, set up omni-channel interfaces - web, mobile, intranet. Assign risk profiles. For example, public systems might allow only deterministic answers. Internal tools may use hybrid mode where internal users are more risk tolerant.

Security is built in. All data stays inside enterprise systems. No third-party LLM training. No leakage. Models are disclosed.

Roles are clear:

- Knowledge Managers approve content
- IT handles integration and security
- SMEs support escalation and accuracy

By launch, your team has oversight. Your AI is a brand-safe operator - not a risky black box.

Future Outlook

80%

OVER 80% SAY AI HASN'T MOVED ENTERPRISE-WIDE EBIT IMPACT. ONLY 17% ATTRIBUTE ≥5% EBIT TO GENAI [2]

MCKINSEY & CO, 2025

As AI evolves from support to execution, governance becomes more and more important. Responsible AI won't just help - it will help orchestrate systems to resolve problems. As kama.ai's Hybrid AI eBook states, "Hybrid AI Agents represent the future of digital labour, where AI not only understands enterprise knowledge but applies it with accountability, transparency, and real-time impact."

AI governance - like explainability, containment, and audit-ability - will become as vital as cybersecurity and identity management in the IT sector. AI will shift from tools to or-chestrators. It will run workflows across HR, legal, finance, and support functions. It will ensure brand tone, legal compliance, and smooth task completion are all in order.

Procurement will also evolve. Buyers will focus on trust metrics: audit logs, hallucination rates, and response traceability. The safer and more responsible the AI, the better its value.

Brands will market AI safety. Just like eco-certifications, responsible AI will drive trust. Public AI will reflect brand integrity.

Zero-trust deployment models will rise. Each AI response will be verified, logged, and scored for safety or accuracy.

The future of artificial intelligence is one in which responsi-ble AI becomes a strategic edge - not just a safeguard. This is especially true when queries have actions associated with real-world consequences and impacts.

Key Takeaways

Found this eBook Useful?

Pass it to a Colleague

Help your team stay ahead with safer, smarter AI. Forward this PDF.

Complex tasks are the next evolution of AI in the enterprise. This is the state in which we will ask our AI virtual assistants to take actions on our behalf. But they will need more than just fluent answers. In this new era, AIs will demand verified action, safe automation, and brand alignment.

Hybrid AI offers this. It separates fact from fiction, answers from actions, and risk from reliability. By combining deterministic and generative tools, these systems can automate work without losing oversight.

The result is a trusted teammate. One that's scalable, safe, and enterprise-ready.

Here's what you need to remember:

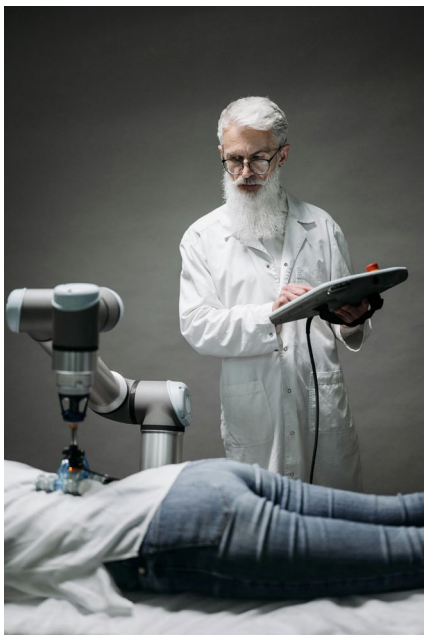

- AI's next wave isn't about chat—it's about doing.
- GenAI alone isn't safe for high-risk or regulated tasks.
- Hybrid AI blends accuracy, automation, & brand safety.
- Enterprises must focus on governance, containment, and human oversight.
- Deploying complex-task AI successfully means planning for trust first.

If you're ready to evolve, Hybrid AI is how you move from questions to outcomes, with confidence.

AI adoption doesn't need to slow down. It needs to smarten up.

Let's Build Responsibly

Book a free consultation at kama.ai to see how Hybrid AI can move your business from uncertain answers to trusted action.

Let's connect to check your options.

kama.ai

E: INQUIRIES@KAMA.AI
M: +1 (416) 275 - 1780
W: KAMA.AI

Ready to move beyond chatbots to AI Actions?
At kama.ai, we help enterprises deploy human-centered Hybrid AI systems that don't just answer - they act. kama.ai's platform blends emotional intelligence, brand-safe automation, and deterministic accuracy into one action oriented solution.

Whether you're automating support, streamlining onboarding, or enabling 24/7 service, kama.ai delivers trusted results, without hallucinations.

Want to build your own Complex Task AI system?
We also consult with organizations ready to design safe, scalable Hybrid AI solutions.

Think kama.ai for:

Trust, Empathy, and Intelligent Actions

Responsible AI

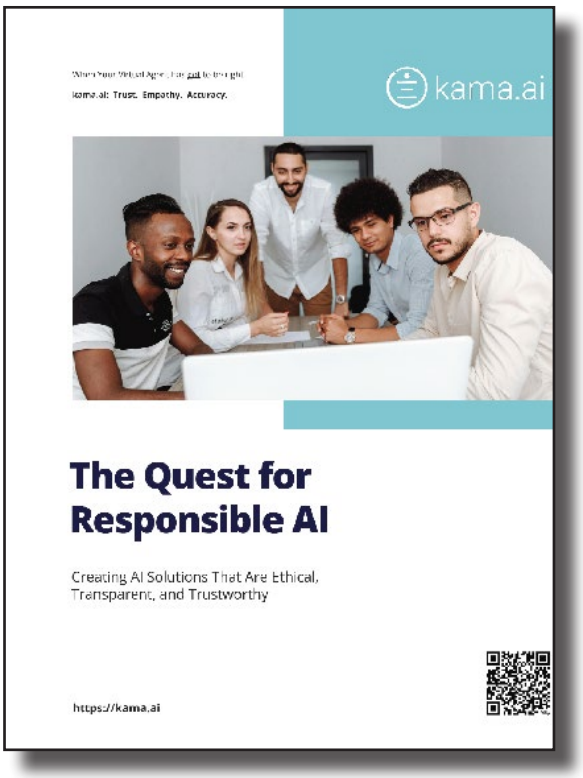

Unlock the future of Virtual Agents with Responsible AI—where trust, ethics, and human values drive smarter, safer, and more meaningful enterprise interactions.

Explore how responsible AI can transform Virtual Agents into trusted, ethical knowledge partners. This ebook outlines how kama.ai's platform combines explainability, human-centered design, and governance to deliver accurate, empathetic, and values-driven information. It emphasizes the importance of transparency, privacy, and environmental sustainability in today's AI landscape.

Discover how kama.ai ensures Virtual Agents align with human values and organizational goals—building trust, protecting data, and connecting authentically with diverse audiences.

FREE to download:
NO STRINGS ATTACHED !

Hybrid AI Agents

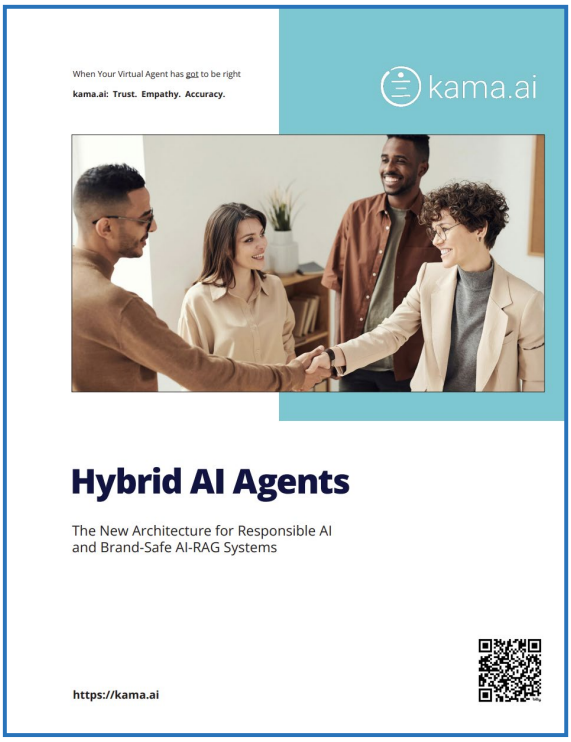

A deep dive into how Hybrid AI Agents are redefining enterprise automation, blending deterministic accuracy with governed GenAI to deliver brand-safe, auditable, and scalable virtual agent solutions.

Discover how Hybrid AI transforms the future of enterprise Virtual Agents.

This ebook outlines kama.ai's approach to Responsible AI through a new architecture that combines deterministic Knowledge Graphs with governed Generative AI using Retrieval-Augmented Generation (RAG). The result is safe, scalable, and brand-aligned AI Agents.

Learn how Hybrid AI eliminates hallucinations, ensures 100% accuracy where required, and lets human experts govern risk. With GenAI's Sober Second Mind®, kama.ai separates creative drafts from factual delivery. This ensures trust, empathy, and accuracy. Trusted Collections, auditability, and continuous human oversight drive enterprise-ready AI that protects reputations and complies with legal and regulatory needs.

FREE to download:
NO STRINGS ATTACHED !

Endnotes

1. **Oxford University**, "Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models", Journal of Legal Analysis, Vol 16. Issue 1. Matthew Dahl, Varun Magesh, Mirac Suzgun, Daniel E. Ho, Ferris Ayar, Jun 2024.

2. **McKinsey & Co.** "Superagency in the workplace: Empowering people to unlock AI's full potential" Jan 2025.

3. **Harvard Business Review**, "AI And Machine Learning: Embracing Gen AI at Work", H.James Wilson and Paul R. Daugherty. September - October 2024.

4. **Deloitte,** "Deloitte Survey Reveals Smart Manufacturing Is Driving Advantage But Needs Focused Investment and Implementation." May 2025.

5. **kama.ai**, "Trustworthy Hybrid AI - The Responsible AI Framework for Brand-Safe Virtual Agents" Feb 2025, https://shorturl.at/62vH9

6. **Gartner,** "Gartner Survey Finds Generative AI Is Now the Most Frequently Deployed AI Solution in Organizations," May 7, 2024.

7. **Deloitte, "**Scaling Agentic AI is a business transformation – not just a tech project" Stu Scotis, July 2025 (posted in The Australian).

8. **Flobotics**, "RPA Reloaded: 50+ Robotic Process Automation Statistics to Know in 2025," Yana Pokora, Content Manager, April 2025.

9. fd

Photography Thanks To:

1. Pexel: Pavel Danilyuk
2. Pexel: ThisIsEngineering

kama.ai

E: INQUIRIES@KAMA.AI
M: +1 (416) 275 - 1780
W: HTTPS://KAMA.AI

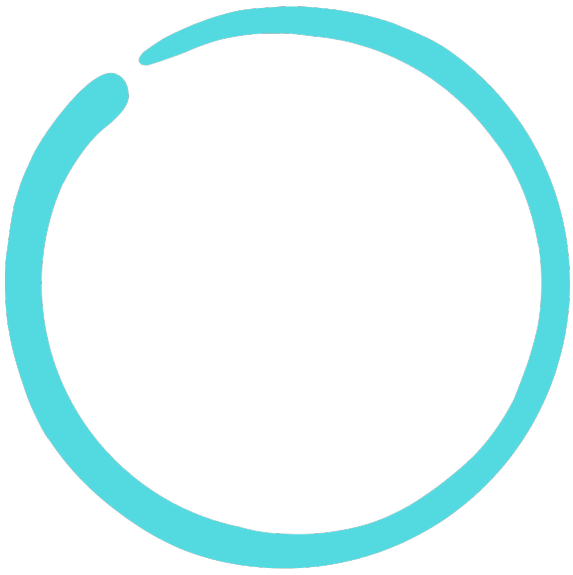

Visit us on **LinkedIn**