omega point.

Moving fast and securely with agentic AI

How to ensure robust cybersecurity, operational resilience and regulatory compliance as you race to transform your business with agentic AI

aws PARTNER Advanced Tier Services

aws PARTNER Security Services Competency

Contents

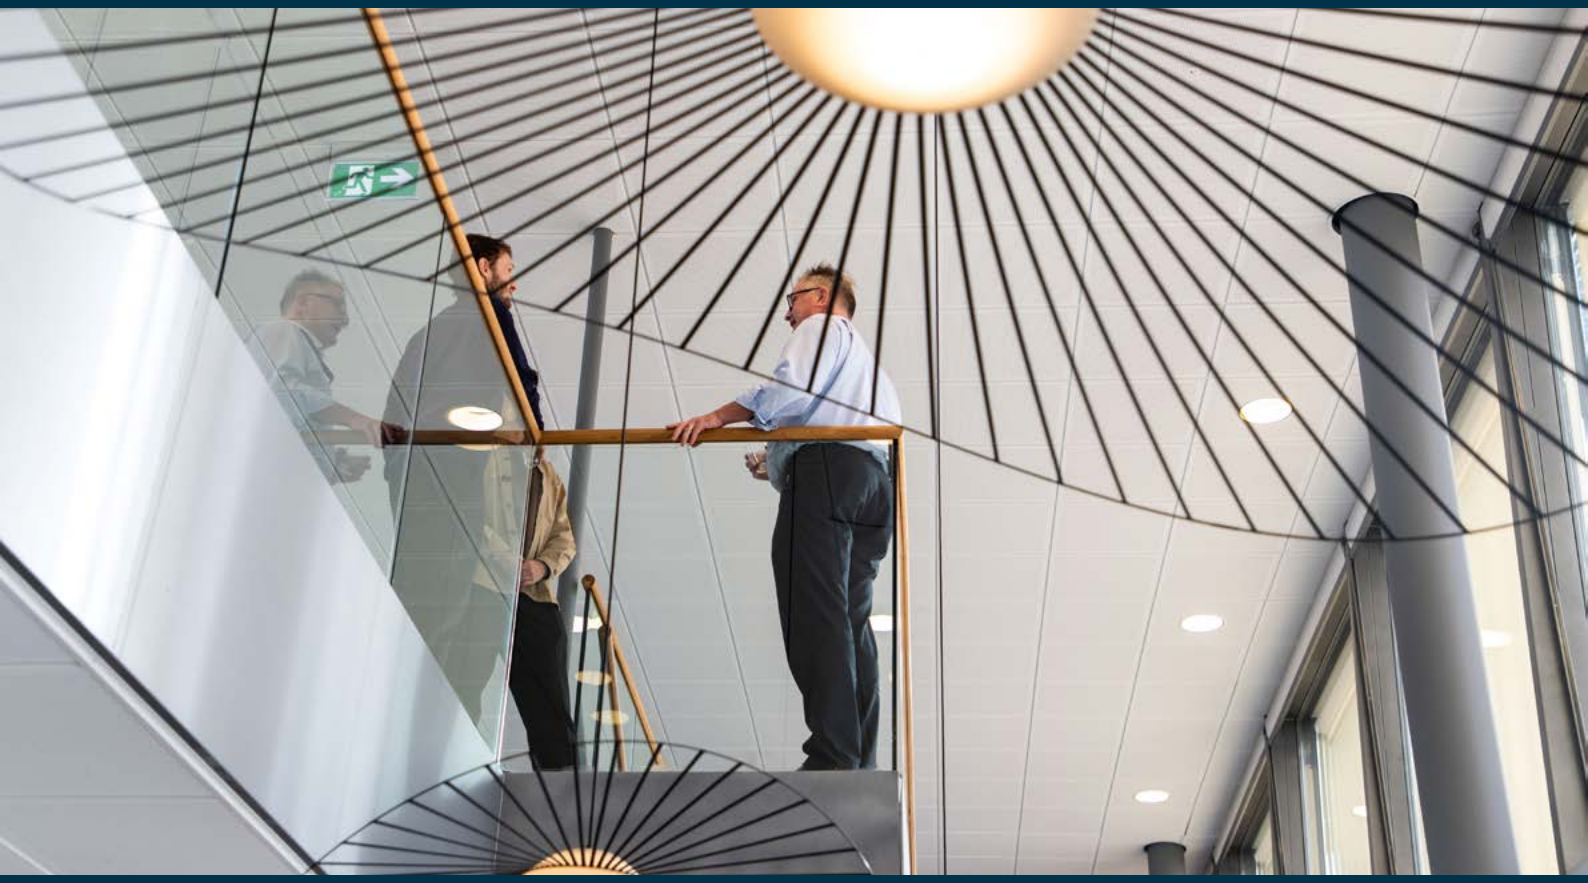

Agentic AI: move fast without breaking things

With business leaders predicting that agentic AI will be as transformative as the creation of the internet itself[1], companies are now racing to integrate this technology into their business models. The race is on to create autonomous AI agents that can perceive, decide, and act across business systems.

In this AI arms race, it's vital to be aware of how new cybersecurity risks may disrupt your operations, lead to regulatory issues, or damage your company's reputation. As AI agents act autonomously, they can intensify challenges to confidentiality, integrity, and availability, amplifying foundational risks, such as data privacy breaches, denial-of-service, and system integrity failures.[2]

These barriers to agentic AI adoption can be overcome. But to do at pace, you must consider AI security "by design". This starts with extending existing threat modelling frameworks to cover autonomous systems and redesigning security and governance guardrails to handle agentic systems.

This whitepaper will help you to build these considerations into your AI strategy, so you can adopt agentic systems securely and at pace. It summarises a strategic approach that will empower you to bring new agentic solutions to market rapidly, safe in the knowledge that they will be secure and compliant by design.

[1] AWS re:Invent 2025, Keynote with CEO Matt Garman, 2025
[2] McKinsey, Deploying agentic AI with safety and security, 2025
[3] Deloitte, The agentic reality check, 2025
[4] Deloitte, The agentic reality check, 2025

Agentic AI will transform business operations

of day-to-day work decisions will one day be made autonomously through agentic AI[3]

of enterprise software will embed agentic agents, up from almost none today[4]

AI security risks you cannot afford to ignore

Agentic AI systems have the power to access and interact with IT systems and data, directly triggering transactions, changing records or initiating workflows. This means any security issues can have immediate business impact—from incorrect payments to corrupted data or halted operations.

As your teams innovate to unlock agentic AI's potential, it's vital that they are aware of the cyber risks unsecure systems may expose your business to, so they can take action to secure their systems.

These potential risks include:

Malicious prompt injection.
Attackers use malicious prompts to bypass agent security protocols. This can allow them to access sensitive data or influence the agent's behaviour.

Data poisoning.
Attackers inject falsified information into the model's training data in order to influence model decision-making and change an agent's outputs.

Tool misuse and exploitation.
Attackers subvert the agent's toolchain or exploit weak validation in the orchestration layer to get the agent to do things that it shouldn't.

Model bias.
All AI systems have inherent biases that reflect biases in the datasets that trained them, as well as those of the engineers who programmed them.

One survey found that **80% of organisations** have already encountered risky behaviours from AI agents, including improper data exposure and unauthorised access, underscoring the need for robust security and ongoing monitoring.[5]

However, these risks can be mitigated with the right cybersecurity frameworks, tools and processes. Enlisting a trusted AWS Consulting Partner, such as Omegapoint, can be an effective way to ensure you have the right risk management processes in place to continue innovating at pace.

[5] SailPoint Technologies, AI agents: The new attack surface, 2025

Considerations for cybersecure agentic AI

DATA PLANE

Model selection
- Model operating model
- Model ownership
- Model isolation/sandboxing

Model development
- Privacy and regulatory controls
- Cybersecure software and CI/CD
- Cybersecure cloud platforms

Model training
- Training bias
- Model drift
- Model provenance

AGENT PLANE

Solution design
- Follow technology platform security recommendations
- Penetration tests
- Threat modelling

Identity access management
- Role-based access control
- Secrets management maturity
- Identity governance

Cybersecurity
- Zero Trust
- Secure by design
- Logging and auditing

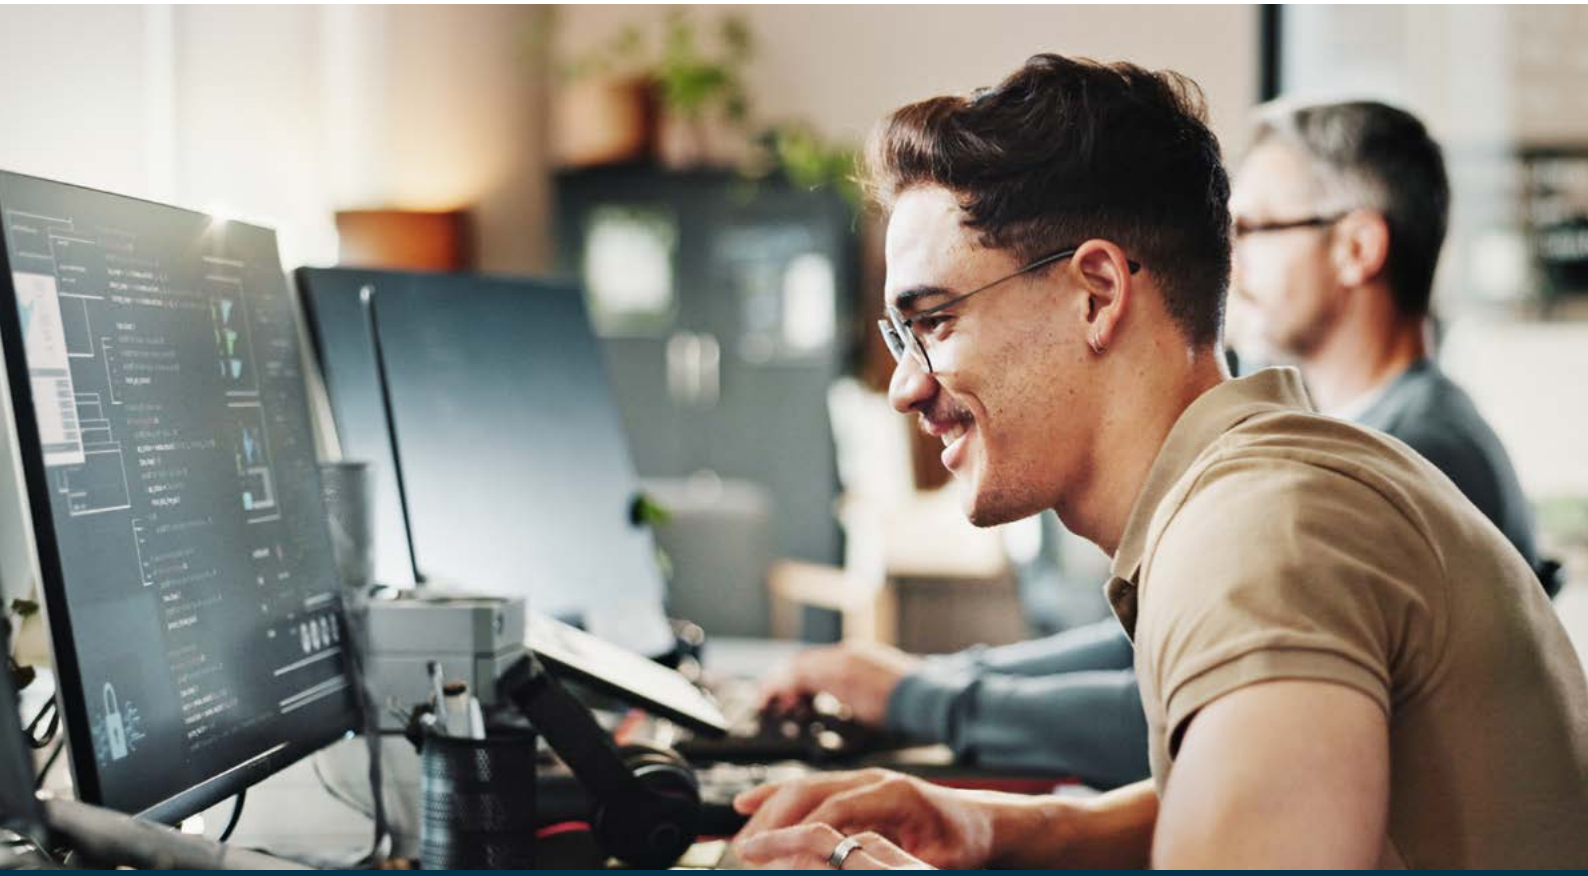

Ensuring your AI agents are secure by design

AI security isn't something you can easily bolt on after your agent has been built. It must be considered from the outset, so you can put processes in place to ensure each project phase is delivered with security and resilience in mind.

Here are some of the key things we consider at Omegapoint when embarking on agentic AI projects for our clients:

1.

Governance and risk ownership

It's important to establish clear joint ownership between the business sponsor, the CISO (or equivalent), and your data/AI leader.

When developing the business case, include a lightweight risk assessment. What data will the agent see and generate? Which systems can it read from and write to? What could go wrong? And what regulatory requirements apply?

2.

Agent-specific threat modelling

Agentic AI introduces new behaviours—autonomy, planning, tool use, persistence—to your IT ecosystem that change your attack surface.

For each agent, your teams should consider what concrete actions it will be able to perform. Which tools and APIs will it be able to call? What operating model will it use? And which abuse and failure scenarios are realistic?

3.

IAM integration

Architect your agents like high-risk services. Give each agent an identity and role in your IAM (identity access management) system, scoping its permissions to the minimum required and enforcing segregation of duties between agents where necessary.

Isolate agents from each other and from core systems via well-designed APIs and network controls, and introduce a dedicated orchestration layer that enforces policy, validates parameters, and applies rate limits. This should also provide consistent logging and auditability across all agent actions.

4.

Observability, guardrails and kill-switches

You need to be sure that, if something does go wrong with your agent, your team will see it quickly and contain it safely. That means ensuring model activity is observable and alerts are triggered whenever a high-risk action occurs. (For example, large transactions or data exports.)

You should also put guardrails in place that prevent foreseeable issues and enable rapid responses in the event of an incident.

5.

Protecting data across the agent lifecycle

A secure-by-design approach requires you to look at training, retrieval and runtime as separate—but connected—risk domains.

Key practices include data minimisation and scoping to ensure agents have access to the minimum data required for their tasks, segmentation of domains to keep sensitive data protected, and policy-aware data retrieval and retention.

Your solutions must enforce access control, purpose limitation, and consent rules in whatever retrieval layer feeds your agents.

6.

Aligning people, processes and platforms

Upskill and enable your teams to give architects, engineers and product owners clear patterns and processes for agentic AI security.

Train operational teams to recognise and escalate suspicious agent behaviour. Integrate agentic AI into existing processes, leverage best practice frameworks for AI security, and choose tools and platforms that offer first-class security features.

Why AWS for agentic AI?

Faster time-to-value
Accelerate from prototype to production with fully managed services that eliminate infrastructure complexity.

Flexibility and interoperability
Build agents your way using any framework, model or tool—while maintaining complete control over how your agents operate and integrate with existing systems.

Security and trust at scale
Deploy with confidence using enterprise-grade security features designed to ensure your agents operate reliably and securely at scale.

Deliver resilient, secure AI agents with Omegapoint

Gain assured resilience against digital attacks without sacrificing speed-of-change or innovation pace with Omegapoint and AWS.

As a trusted AWS Consulting Partner and Northern Europe's leading tech company in cybersecure digitalisation, we have 25 years' experience helping companies across the globe develop applications and AI solutions that are secure by design. We'll apply our proven "secure by design" methodology to help you deliver secure, resilient agentic solutions at pace.

STEP 1

AI feasibility study

We'll help you understand which AI initiatives are technically and commercially realistic for you. The goal is not to sell visions, but to provide a high-quality decision basis founded on the actual maturity of the technology and the customer's strategy.

We use a clear maturity ladder to distinguish between established applications (RAG, semantic search, summaries), more exploratory projects (generative features), and speculative ideas that are not yet mature. In this way, management, business, and technology can speak the same language about risk, value, and feasibility.

STEP 2

Execution and project delivery

After the feasibility study, we'll help you prioritise the value-adding AI use cases we identify together and take them forwards from proof-of-concept to pilot projects to full implementation to achieve robust, secure, and business-driven realisation. This avoids large, high-risk initiatives and delivers concrete results quickly.

You'll get a scalable delivery team with mixed competencies and seniority levels to ensure a fast, secure, and cost-effective delivery.

STEP 3

Penetration testing

We also provide rigorous penetration testing to ensure your LLMs, agentic applications and AI infrastructure are resilience against bad actors.

Our team can also review your architecture, system design, source code, and processes for secure development and operations, upskilling your teams to adopt best practices for efficient, resilient AI development.

We'll help you apply our "secure by design" principles to maintain an adapted security baseline for AI development and operations.

Why Omegapoint?

- AI solution development that's secure by design
- Security as a core principle, built into every layer from day one
- 20+ years of experience of end-to-end delivery
- Security integrated into architecture, code and processes from the start
- Cloud native—robust, modern and ready for future environments
- Transparency and trust—quality and security at every step
- 300+ applications developed for clients across multiple industries
- Business-critical systems—web-based administrative solutions, tightly integrated with processes with reusable modules—built efficiently on proven components and battle-tested methodology
- Long-term partnerships built on satisfied clients, on-scope deliveries and clear cost management

Unleashing efficiency gains with AI chatbots

We're already helping businesses transform securely at pace

The challenge
HVAC (Heating, Ventilation and Air Conditioning) systems manufacturer and wholesaler Purmo Group's processes for handling its product catalogue were cumbersome and error prone.

This led to slower go-to-market, significant time spent troubleshooting, and loss of revenue from products not being displayed in digital channels.

The company is present across all of Europe, with 12 separate brands and almost 500,000 stock-keeping units, where each unit can be sold under many different brands in the same market.

Our solution
Omegapoint extended the capabilities of the company's Product Information Management (PIM) solution, Inriver, with an AI-based chatbot that can diagnose issues, provide remediation suggestions, and automatically adjust the catalogue according to those recommendations.

We co-ordinated and led the development of this solution, providing expert consulting with competencies in solution architecture, ERP and WMS, Inriver PIM, and AI/LLM chatbots.

Results
The chatbot resolved 100% of the questions it was asked in the two months following its deployment. This freed up time for employees to focus on more strategic activities by eliminating the need to consult a PIM system expert.

Our solution also accelerated go-to-market of new products by automatically aggregating, validating, enriching and synchronising product data across channels in real time. This minimises manual work, errors and iteration cycles before launch.

Securely accelerate your agentic AI journey

Contact our team to discuss your needs and put your business on the fast-track to agentic AI solutions that are secure, compliant and resilient by design

Get in touch

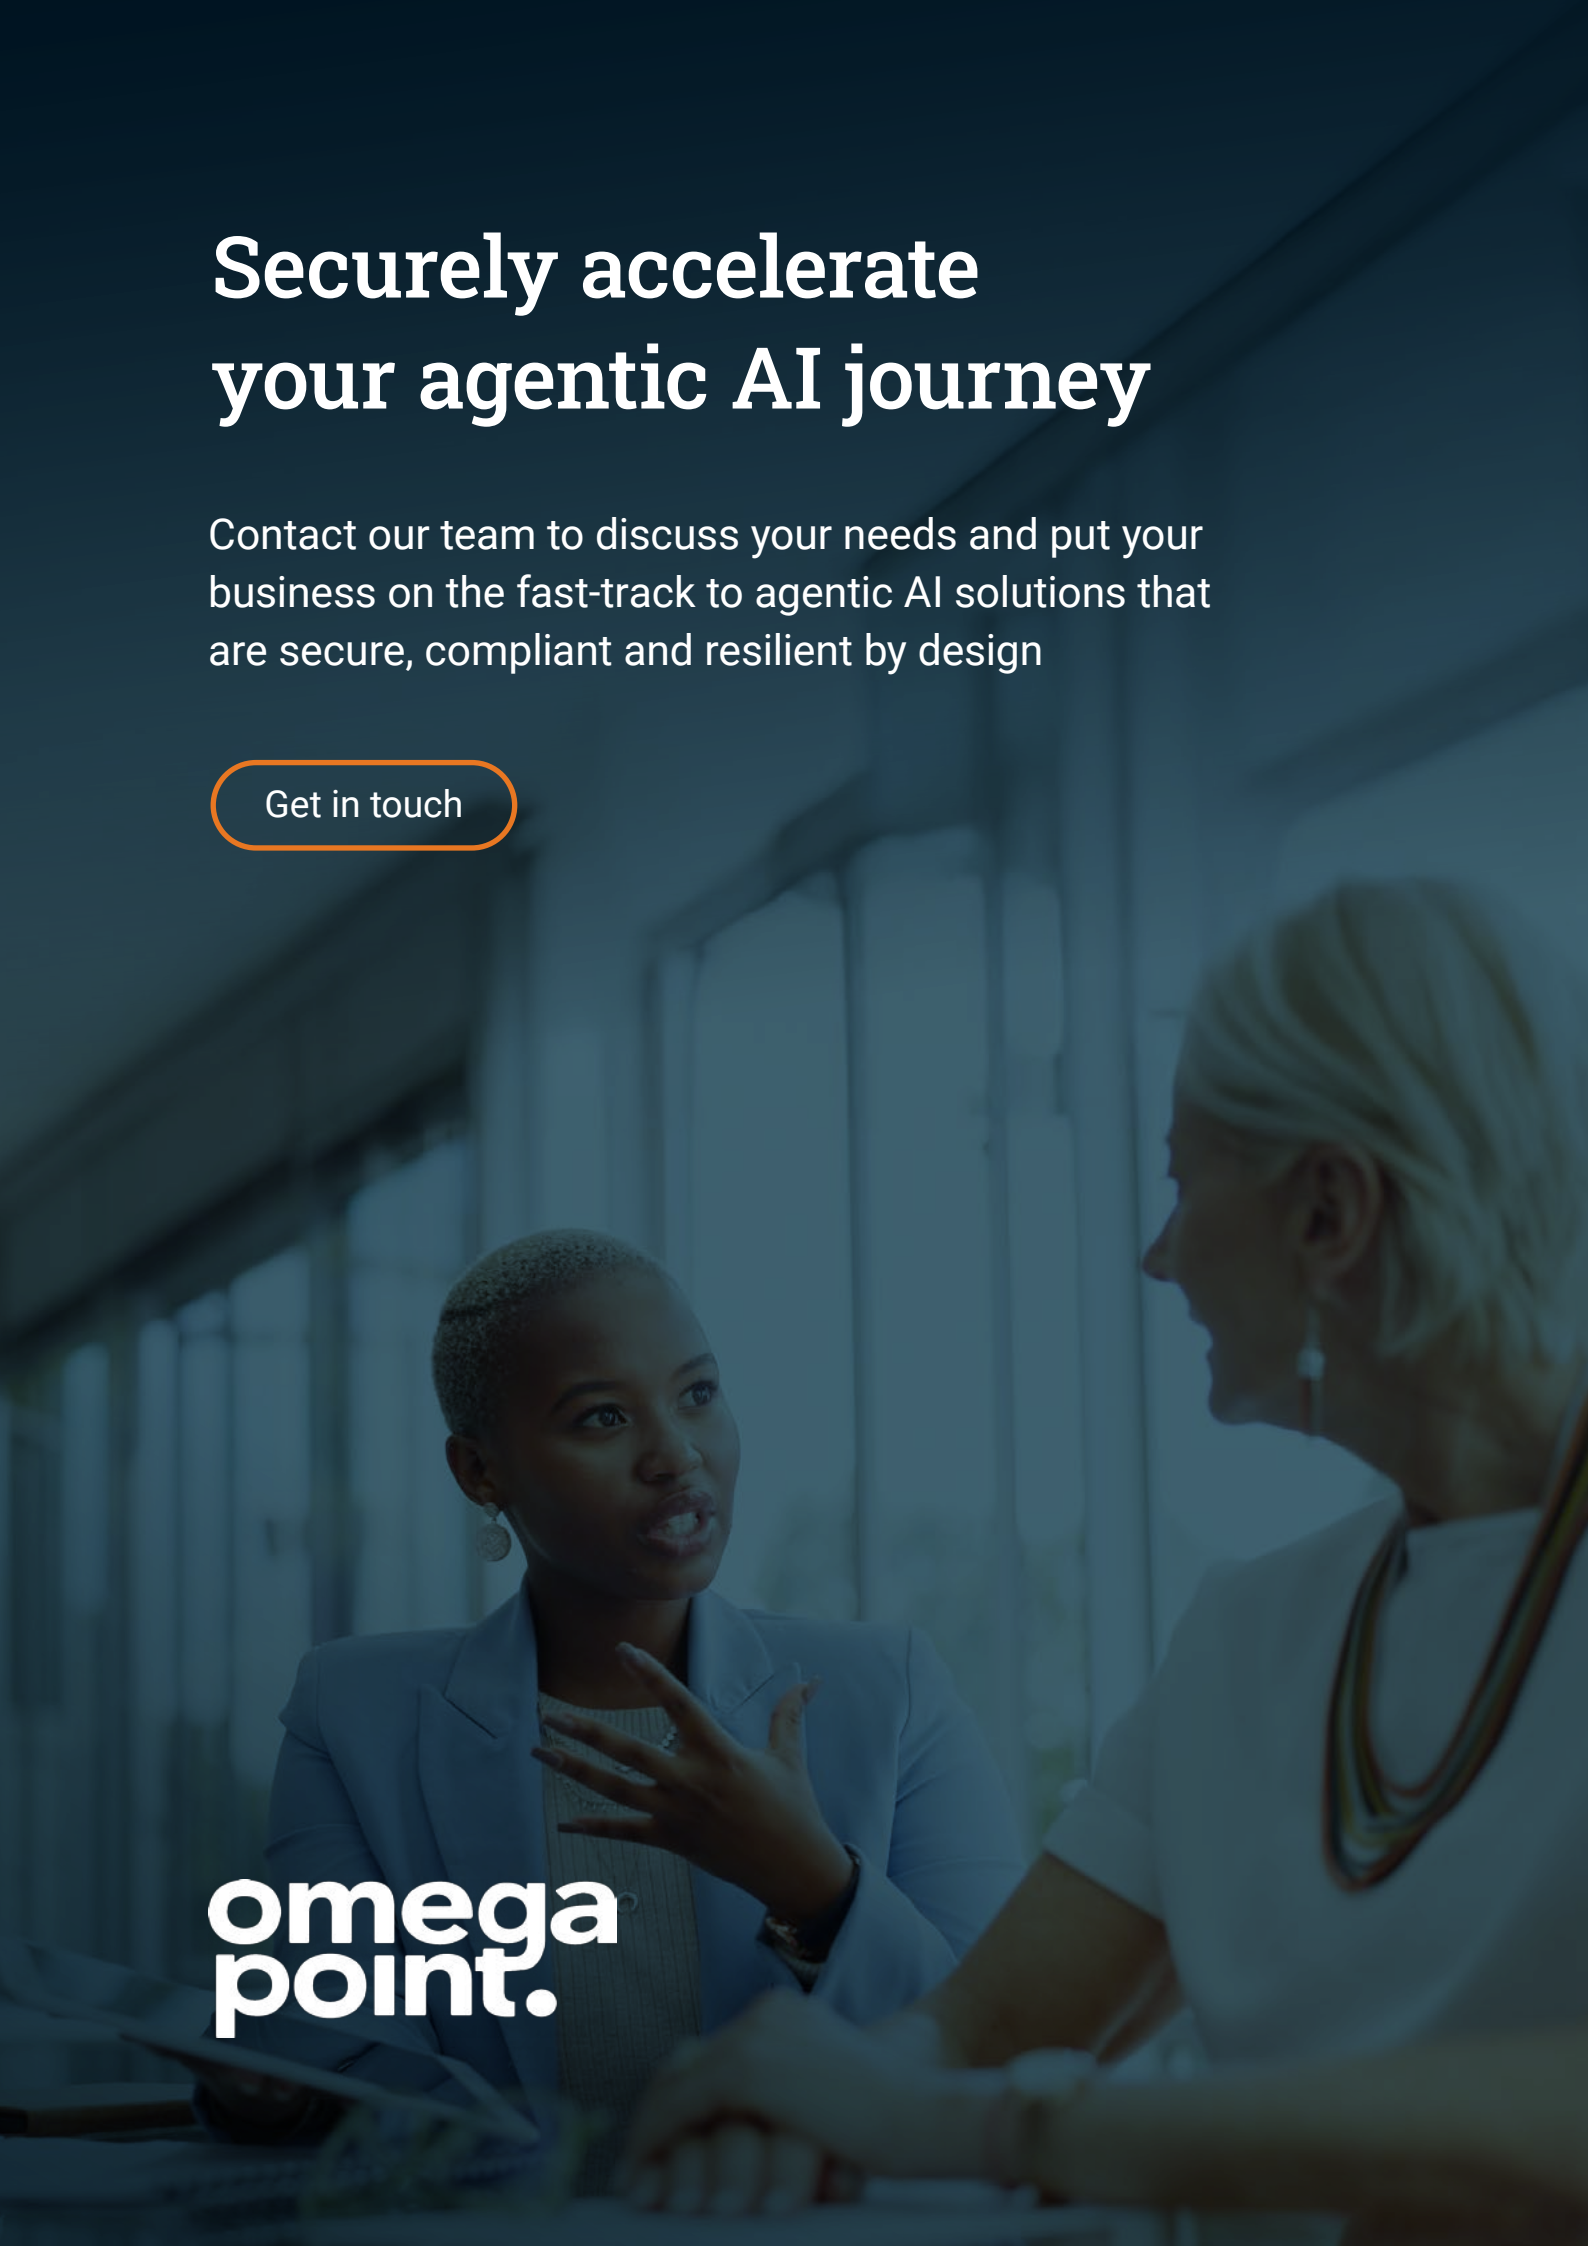

omega point.