CIKLUM

Agentic AI at Scale

A Decision-Maker's Guide to Sustainable Implementation

James Lennon,

VP Experience Engineering

James Lennon is a technology trailblazer transforming software development and user experiences through AI. As VP of Experience Engineering at Ciklum, he drives the adoption of AI systems to accelerate delivery and exceed customer expectations. With over 15 years driving digital innovation, James combines deep hands-on experience with a vision for how AI can overhaul software design, development, and testing workflows. James leads a global engineering and R&D team shaping the future of how technology defines the way we live and work.

Index

Introduction

Why Agentic AI Is Now a Business Imperative

Modern agentic AI has massively expanded the capabilities of artificial intelligence across organizations and enterprises. These AI systems provide a level of autonomy far beyond that of basic AI applications, which rely on simple API connections. Agentic AI can make its own decisions, execute tasks, solve complex problems, and interact with other systems and applications.

On a practical level, organizations that have already implemented agentic AI have found that it can:

- Complete tasks based around unstructured data such as deep-level research
- Solve complex problems through interactions with multiple systems and data sources with minimal human guidance
- Remove human bottlenecks and turn workflows that take days into automated sequences that take minutes
- Enable transformative opportunities for personalized customer interaction and service delivery
- Positively impact business valuation and market perception, thanks to all of the above

According to Gartner, one-third of enterprise software applications will include agentic AI in 2028, enabling 15% of day-to-day work decisions to be made autonomously. This means that now is the time to invest in agentic AI if your organization hasn’t done so already.

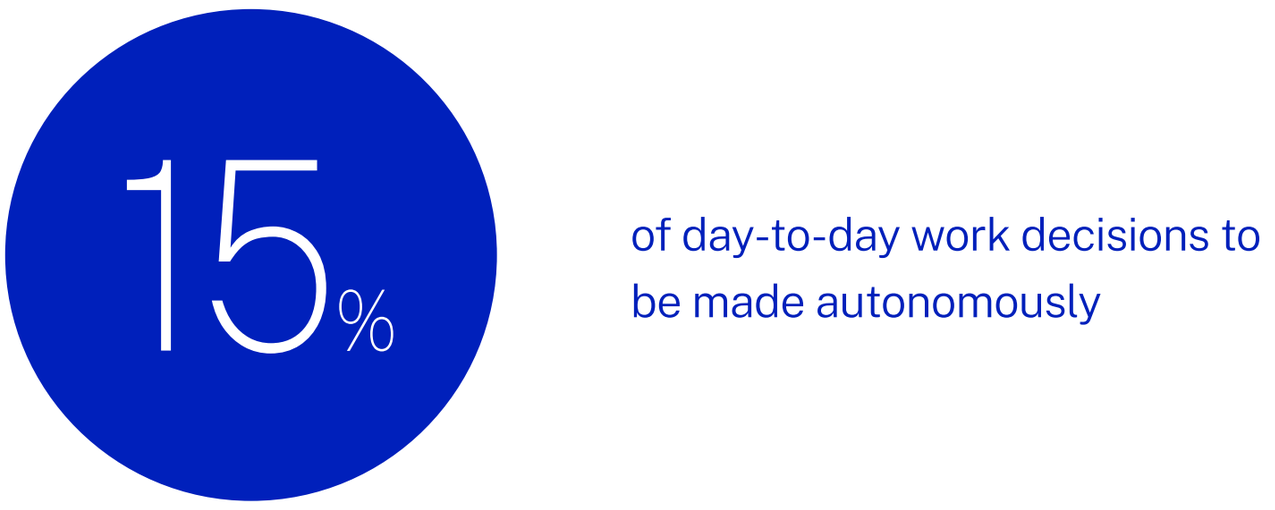

However, care must be taken to get implementation and scaling right; early adopters have struggled when they haven’t built in the flexibility to adapt when underlying AI models are updated or discontinued. Without proper implementation, your organization risks significant costs that can impact return on investment.

This guide explores how a sustainable approach to agentic AI demands a balance of immediate value and long-term flexibility. We’ll highlight the key scaling challenges to address and a strategic framework for tackling them; and how we at Ciklum have developed an effective implementation strategy for agentic AI.

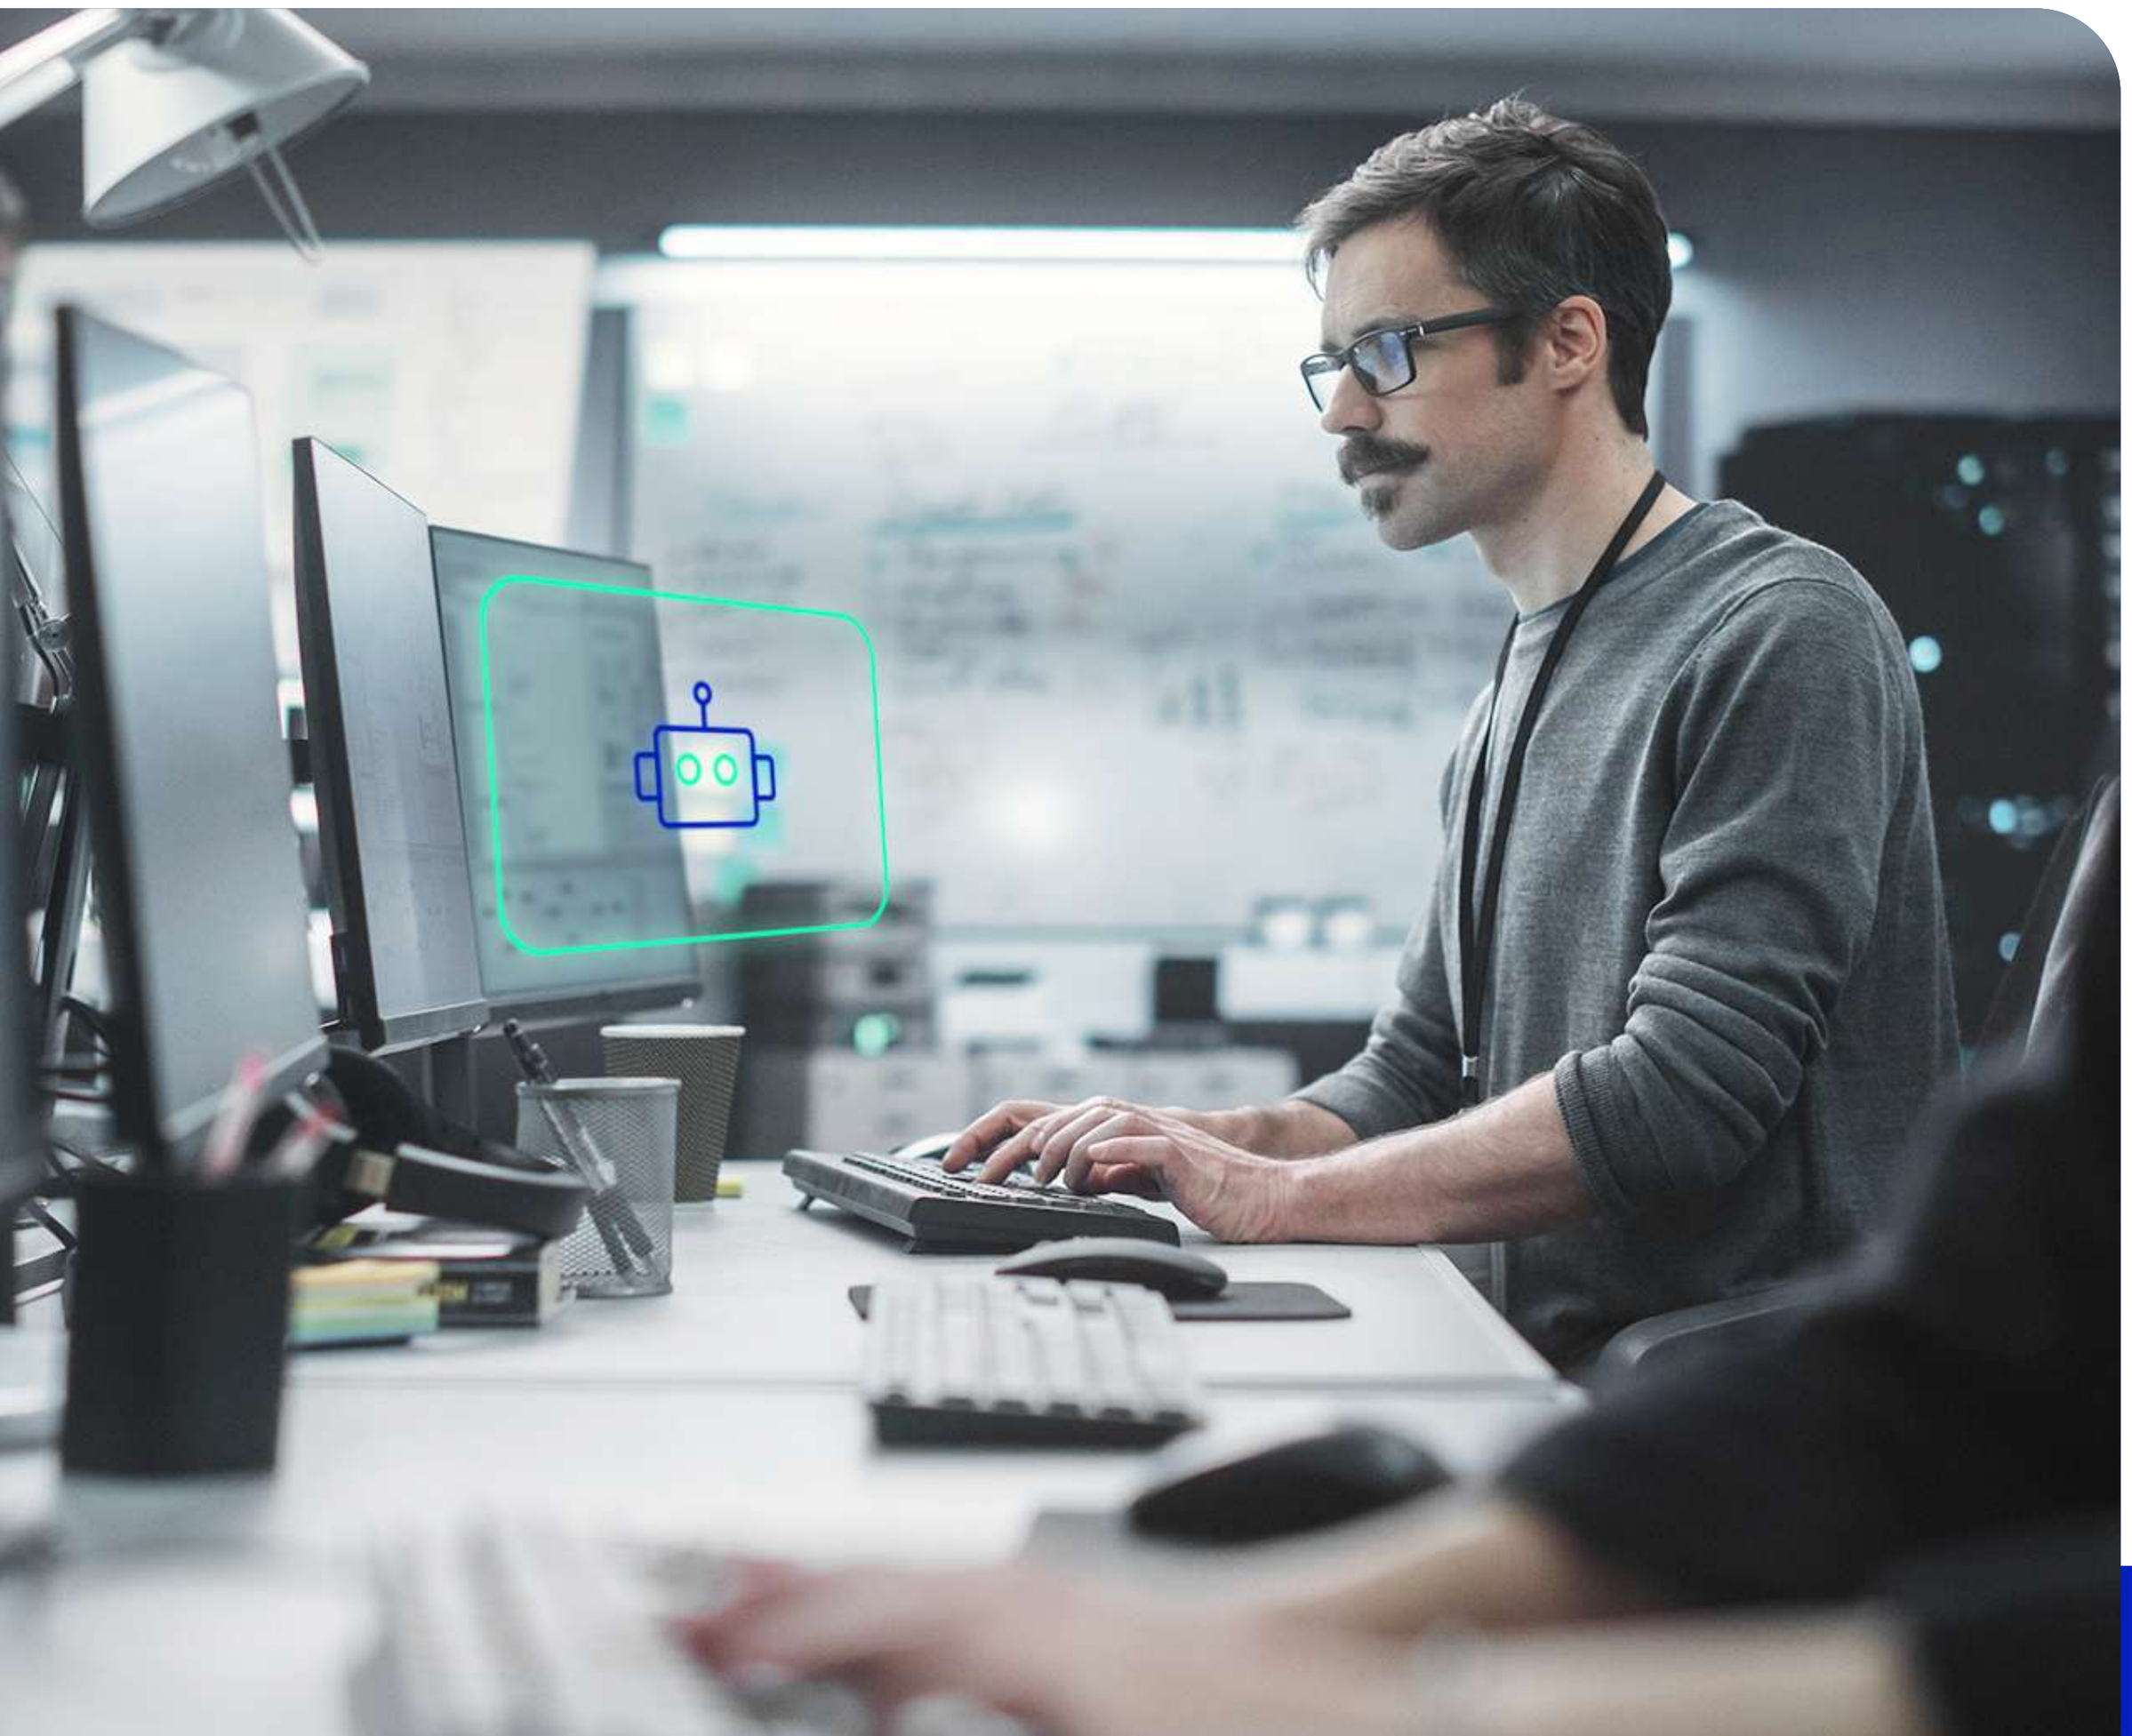

What Stands in the Way of Successful Agentic AI Scaling?

From our experience assisting organizations with their agentic AI implementation, there are four common challenges that should be considered and addressed from the outset:

1 Building in Flexibility

Businesses that rushed to adopt AI as early as possible often did so with the assumption that the underlying models wouldn't change. This proved not to be the case: for example, when GPT-3.5 was discontinued by OpenAI, many organizations had to rebuild their solutions from scratch.

In this situation, organizations are forced to commit significant investment to adapt to new technologies, especially with AI evolving so quickly. And this spending will be unavoidable, as end-user perceptions of an outdated AI solution will be negative compared to the more innovative solutions of competitors.

2 Balancing Autonomy with Oversight

We've found that many organizations find it difficult to determine the right level of autonomy and agency for different use cases. Too little human oversight can lead to security and reliability risks when agents have access to multiple systems, as well as hallucinations (when AI delivers false or misleading output) going unchecked. But too much oversight can lead to human bottlenecks (such as process approvals) that limit the speed and efficiency benefits of agentic AI.

3 Building the Foundations of Enablement and Change Management

The human impact that agentic AI implementation can generate should not be overlooked. Basic AI education is often still lacking, even among AI implementation teams, and there can be cultural resistance irrespective of whether an AI solution is successful or not.

This often leads teams to revert to more familiar non-AI tools, which can be avoided with proper enablement and change management planning. This should incorporate both technical training and "hearts and minds" approaches to drive adoption across the workforce.

4 Choosing the Right Use Cases and Architecture

With AI so popular and attracting so much interest, it's easy for organizations to deploy AI solutions that don't address specific user needs, therefore inhibiting return on investment. The same applies to proper architecting of solutions, and getting this wrong can have significant implications for scalability, security, and adaptability.

It's vital that organizations can distinguish use cases between those where agentic AI is ideal, and those where traditional automation or Robotic Process Automation (RPA) may be more appropriate.

A Strategic Framework for Scaling Agentic AI

Technical expertise is only one key part of making an agentic AI implementation a success. It also demands coordination across multiple departments and functions, so that the right technology is applied in the right areas, and so that any potential weaknesses are addressed in a timely manner. For this reason, a comprehensive, structured framework should be established at the outset, that treats technology as an equal to two other considerations: people and processes.

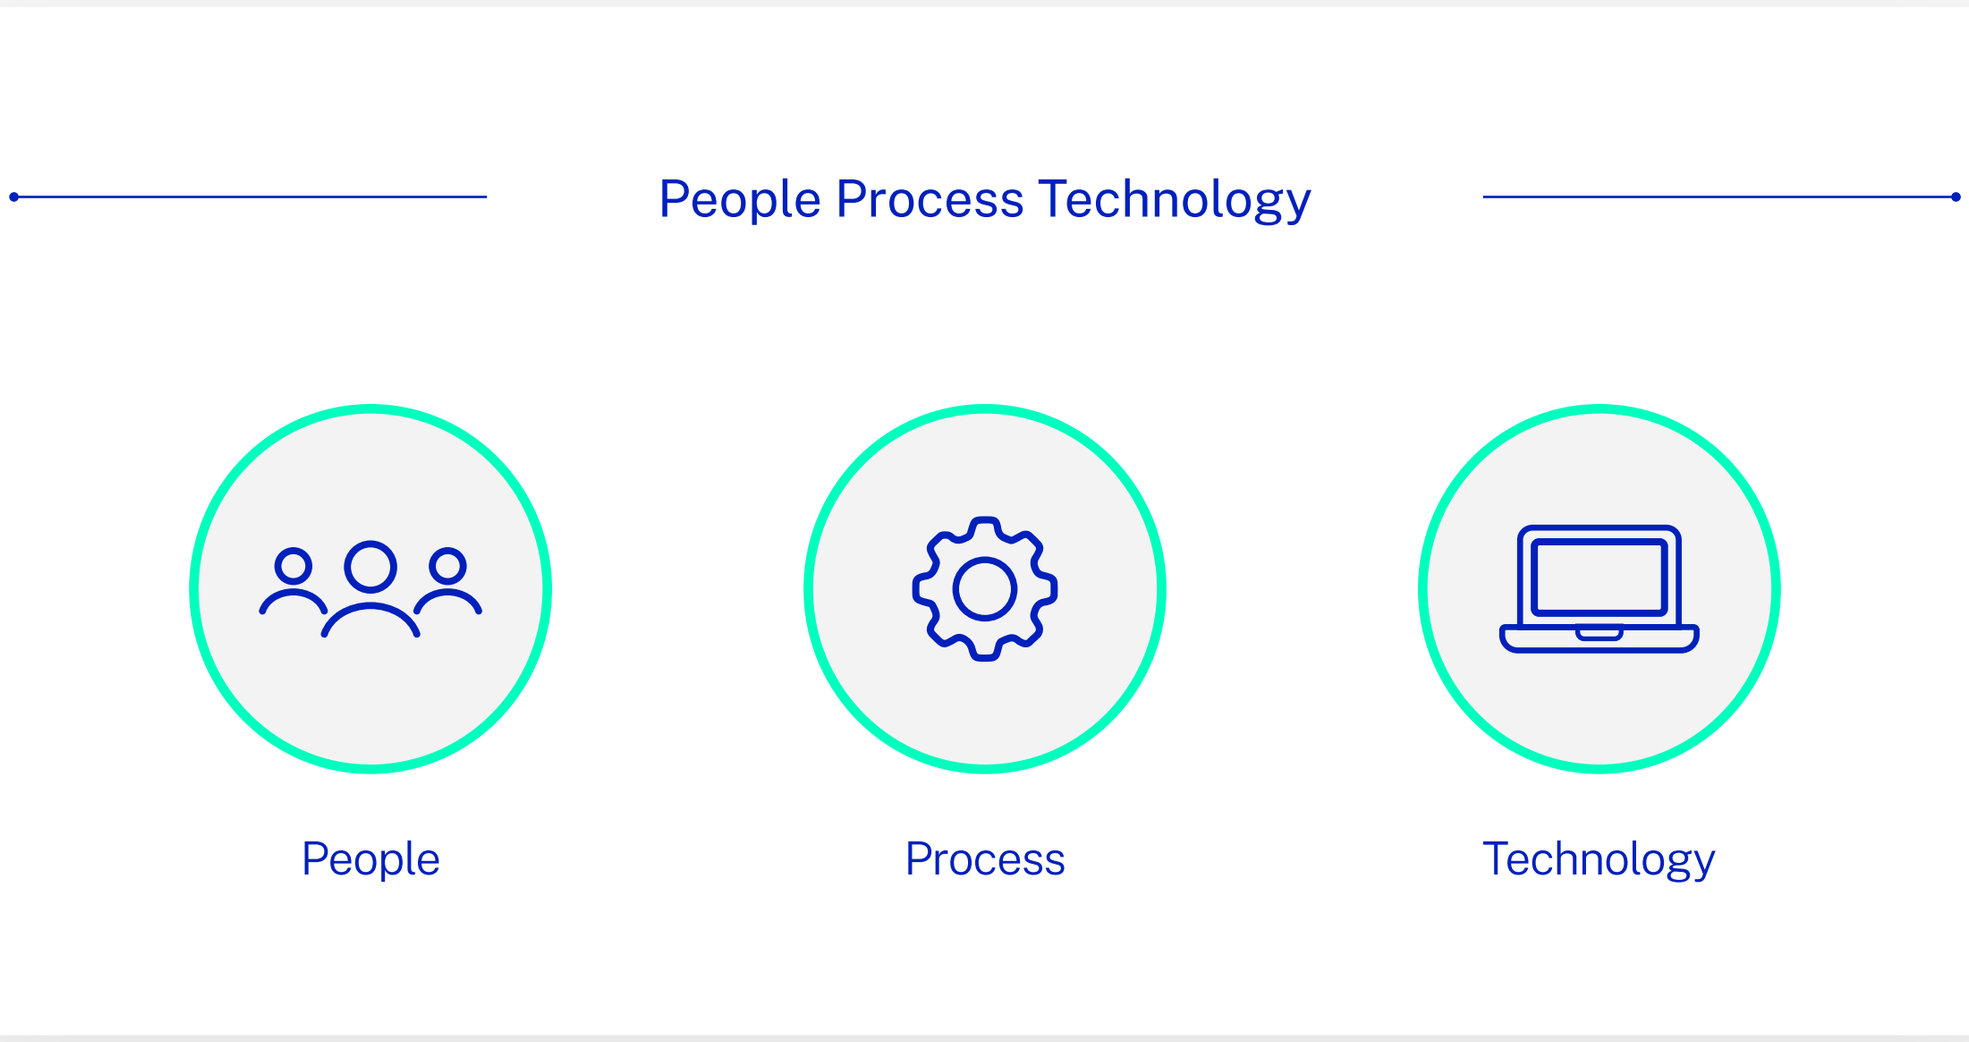

01 People: Enablement, Education and Organizational Readiness

A holistic approach to preparing the workforce for AI implementation should include:

- Basic AI education for everyone involved
- Training on AI system capabilities and prompt engineering
- Sharing of success stories across teams to encourage buy-in
- Promoting self-help resources like Chrome extension hubs so that employees can easily access guidance on AI terminology and applications

02 Process: Identifying High-Value Use Cases

To ensure AI initiatives deliver real value, a product-centric approach to use case selection is essential. This treats AI implementations as ‘products’ that solve specific user problems and generate value accordingly. This means focusing on user needs and pain points rather than the capabilities of the AI tool, applying established product development methodologies, and using standard product metrics to measure success.

In our experience, the best way to evaluate and prioritize AI initiatives is through the PRIME framework (Potential, Reach, Impact, Measurable, Effort/Feasibility, Risk). This is an adapted version of the RICE framework that incorporates AI-specific factors like hallucinations. It helps identify opportunities for immediate value, flags up future opportunities that may be suitable as and when AI technology matures, and prevents AI deployments being created that won’t be used sufficiently.

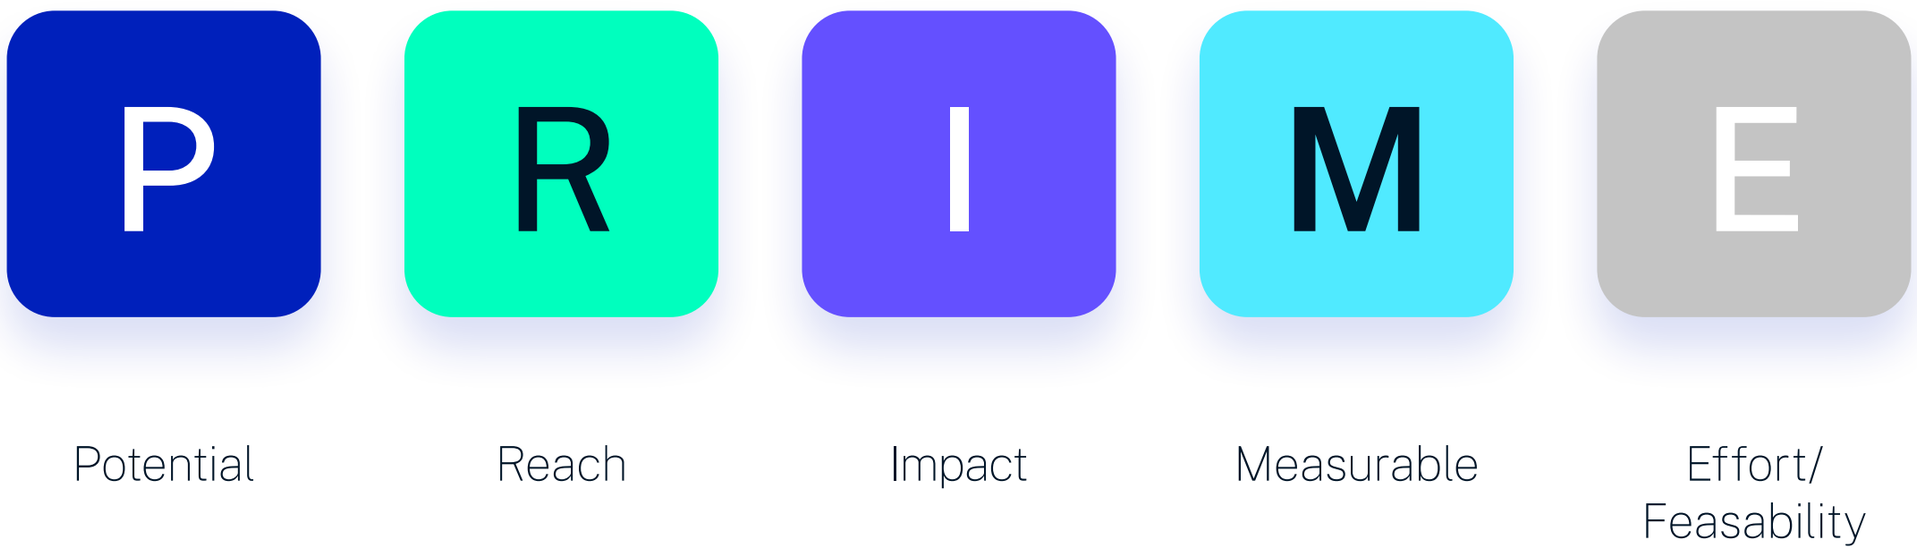

03 Technology: Flexible Infrastructure That’s AI Future-Ready

As early adopters’ struggles have proved, it’s essential to build AI infrastructure that can adapt to whatever the future holds for the technology. This means considering regulatory requirements, organizational culture and empowerment, your desired pace of innovation, and your wider needs around business models and customer engagement.

This involves choosing between three primary architectural approaches:

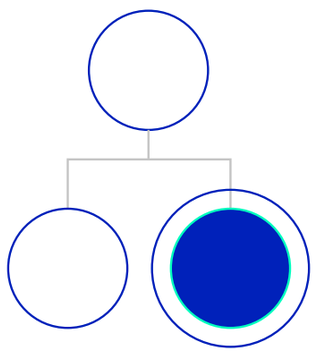

Centralized

Suitable for regulated industries.

High levels of security and standardization, but harder to scale.

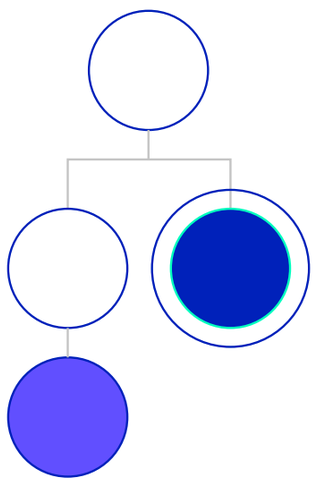

Mesh

Suitable for multi-step information processing.

Balances local expertise with central coordination.

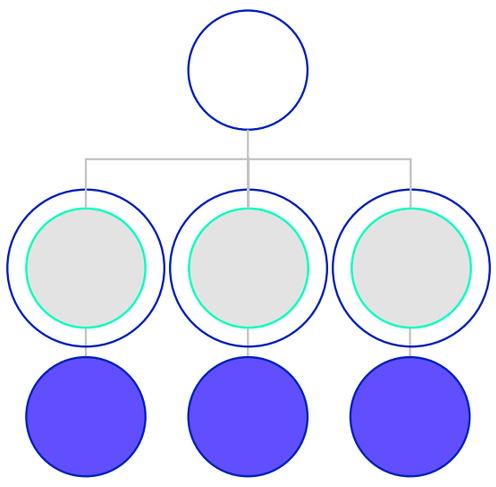

Decentralized

Suitable for tech and product companies.

Enables faster innovation, but can pose global compliance challenges.

04 The Shift from Back-Office to Customer Experience

Much of the focus of agentic AI has been on transforming back-office experiences behind the scenes. But, as this graphic demonstrates, its influence is spreading towards customer-facing functions, and stands to transform customer experiences in the months and years ahead:

Agentic AI has already transformed internal operations and support functions.

- Customer service
  Automated ticket resolution
- Software development
  Bug detection and fixing
- Resolving routine issues
  Without human intervention

Younger generations now expect AI-first experiences.

- Millennials use AI for research and recommendations instead of Google
- Generation Z turn to AI for major life decisions from home buying to career moves
- Consumers now rely on AI as their primary digital assistant

Now is the time for agentic AI to redefine customer interactions.

Personal assistance
‘Take me home’ functions that book rides without opening apps

Service integration
AI handling all interactions with brand apps on user’s behalf

Shopping support
Automated, personalized product research and purchasing suggestions

What are the competitive challenges for retailers?

Will your products consistently appear in AI recommendations?

If customers never directly visit your channels, will your relationships suffer?

Will your brand be commoditized and lose its unique identity?

The answer is a flexible, scalable strategy for agentic AI.

- Optimize for AI discovery by structuring data and content for easy AI detection
- Maintain experience excellence with detailed content and visuals
- Enable agent-to-agent interaction across systems
- Build flexible infrastructure that can quickly adapt to changing trends
- Preserve brand value and proposition even in the age of AI filtering

The Ciklum Way

Our Implementation Strategy

As this guide demonstrates, it's just as important to have a structured implementation strategy in place as it is to understand the theoretical advantages of agentic AI. At Ciklum, we believe a well-planned, sequential approach can address challenges around flexibility, agency and organizational readiness, and strikes the right balance of short-term wins and long-term sustainability that can underpin a successful deployment.

Our implementation approach has four stages, which collectively put the building blocks in place to deliver sustainable AI value. This approach enables organizations to demonstrate AI value quickly while building the infrastructure and capabilities for long-term success:

1. ### Ideation and Proof-of-Concept
   Gathering potential use cases, identifying potential new ones, and prioritizing them by value using the PRIME framework.

2. ### Implementation and Strategy
   Identifying the best technical approach, and building the smallest viable version possible, in order to generate user feedback and demonstrate value quickly

3. ### MVP to Scale
   Delivering solutions through an agile approach, and gaining insights on user feedback and model performance to inform continuous iteration.

4. ### Center of Excellence Capability
   Building repeatable AI capabilities that can validate and process other use cases, and developing a pipeline for future opportunities.

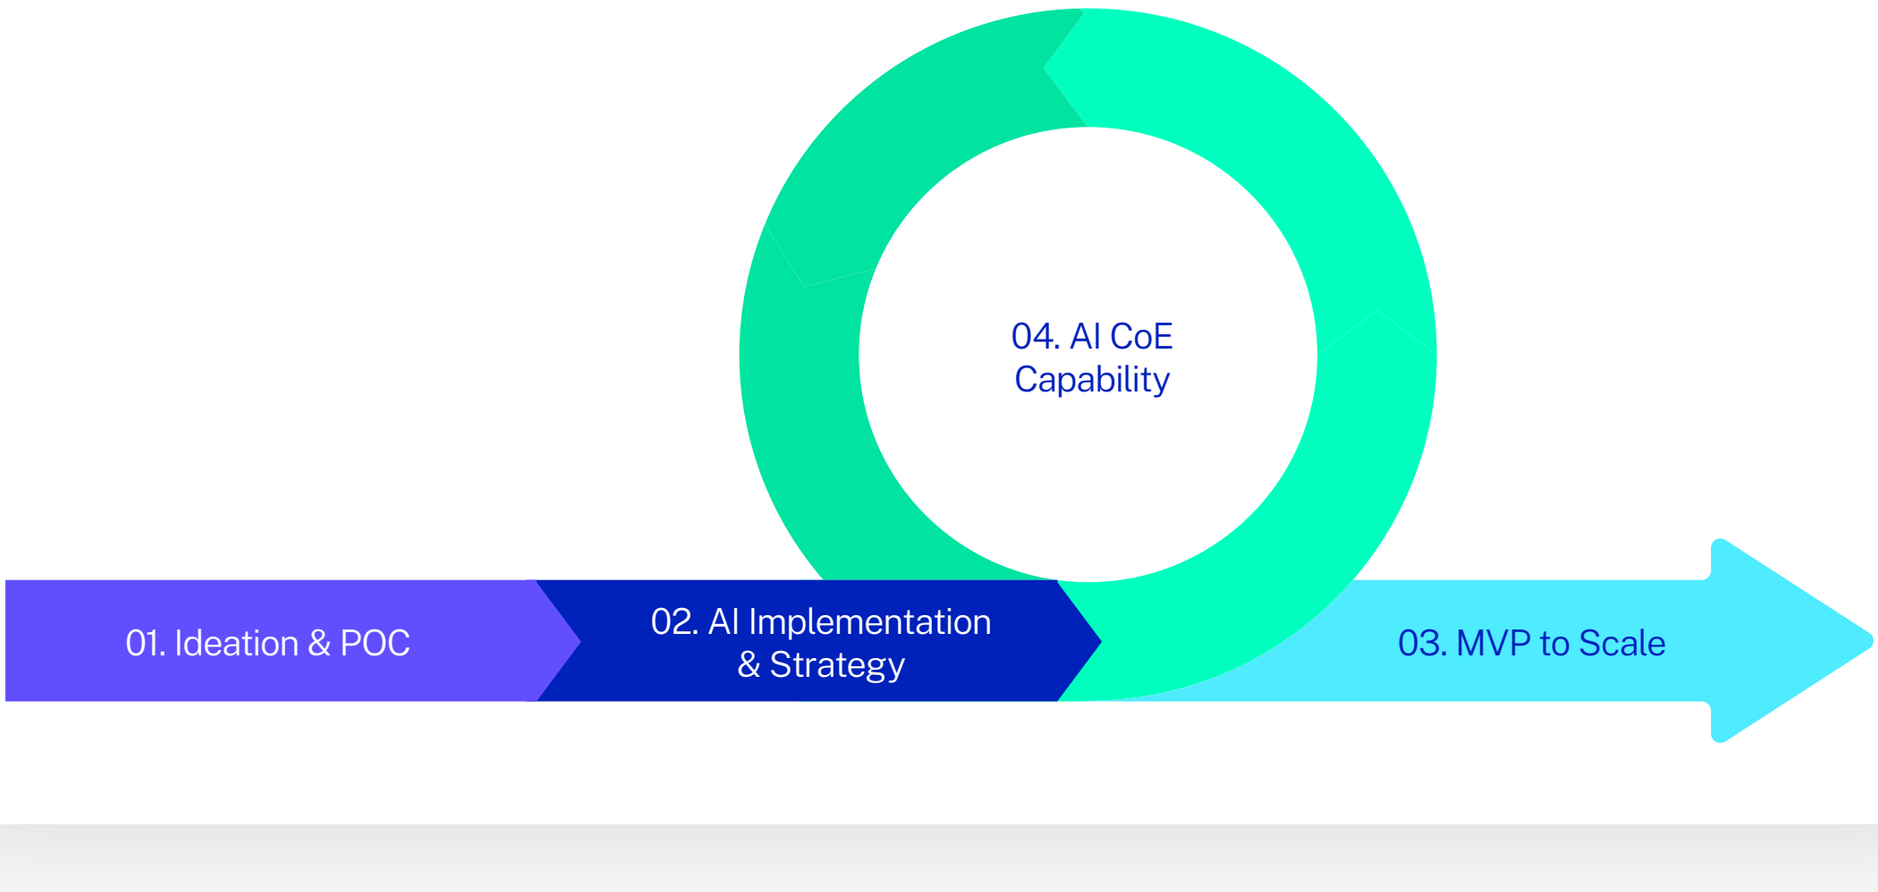

Alongside this process, we incorporate a deep understanding of emerging AI methods and forward-looking expertise. This applies in particular to Model Context Protocol (MCP), which standardizes the context applications provided to Large Language Models; and Agent2Agent (A2A), which enables AI agents to communicate with each other across enterprise platforms.

We believe that, with these technologies set to become industry standards in the months and years ahead, a strong understanding of these emerging protocols is essential for enabling future-proof, adaptable AI implementations.

Putting Our Approach into Practice

We have already been able to successfully apply this approach for many organizations across industries and sectors. These include a multinational online payments company, and a major automotive manufacturer.

Implemented Customer Service Agents for a Global Payments Leader

We leveraged agentic AI to resolve customer service issues and minimize the need for human intervention in the process. The agents can analyze customer problems; search historical data in CRM systems; and identify potential solutions.

When the issues involved are lower-risk, such as unfreezing credit cards after suspicious transactions, the agents can implement fixes completely autonomously. When issues are more complex, the agents can extract tickets from CRM into development systems; write test scripts to try and replicate issues; develop and test code fixes; and then push those fixes into testing - all without human interaction.

These agents have reduced the time demand of customer service on the workforce by as much as 80%, and has enabled our client to reduce its FTE headcount by half.

Enabled Field Technician Support for a Global Automotive Leader

We deployed agentic AI to support the client’s technicians in the field.
The agents act as virtual assistants that help diagnose issues and provide practical solutions. This is based on being able to access decades of specialized vehicle knowledge and analyze that data to come up with the most suitable solution for each specific situation.

This has opened up a vast bank of applicable knowledge to all technicians, regardless of experience. Their SLA has been reduced from six working days to less than one hour, and global support can now be made available 24/7, rather than just during working hours in the United Kingdom. All this has led to a cost saving of over £1 million.

Conclusion

Taking Action for Competitive Advantage

Agentic AI will soon become an everyday technology. Gartner has found that, thanks to agentic AI, as many as 80% of common customer services issues will be resolved entirely autonomously by 2029.

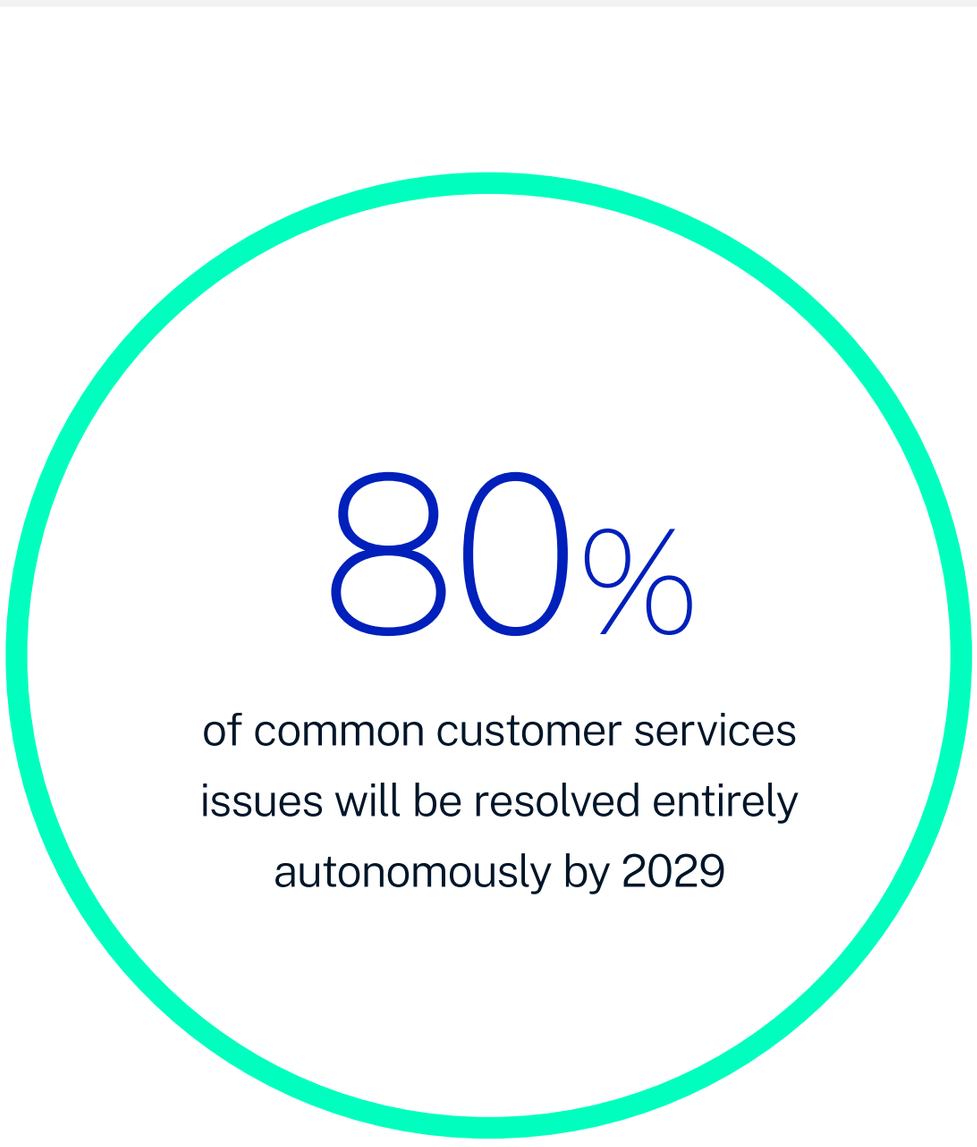

As a result, it is absolutely critical to get any agentic AI implementation right, through a combination of:

- Balancing immediate value with long-term flexibility
- Using a product-centric approach to identify the use cases with the highest value
- Making architectural decisions based on organizational culture, risk appetite, and long-term goals
- Focusing on enablement and change management from the outset

If you can get this right, then your organization stands to realize significant cost savings through automating complex tasks; deliver faster and more accurate service that boosts customer experience and perception; and build internal capabilities to adapt to new AI technologies and more AI-centric customer relationships.

Working with an expert partner like Ciklum maximizes your chances of success. We can assess your organizational readiness, help you identify high-value use cases, and use these insights to select the ideal architectural approach for your needs. From there, we implement targeted solutions that quickly demonstrate value, build on that success through continuous evaluation, and develop a comprehensive AI Center of Excellence.

To find out more, explore our AI-powered product engineering services here, and contact the Ciklum team to discuss your specific requirements.

Contact us today