TERCERA

WHERE SERVICES CAN & SHOULD PLAY IN THE AI ERA

Contents

01 / The AI Paradox for Services

AI represents the single greatest opportunity and the single greatest disruption that services firms will face in the next two to five years.

THE OPPORTUNITY

Gartner estimates that the AI services opportunity will reach $1 trillion by 2030. For context, that's bigger than the GDP of Switzerland and the current U.S. defense budget, in less than four years.

Gartner breaks this opportunity down into two somewhat equal categories of spend:

- Direct: Where AI plays a central role in client projects and is the key focus
- Indirect: Where AI plays a supporting role

That $500 billion of direct AI services spend is broken across Advisory Services, Application Services, and Infrastructure Services, all of which are growing in the mid-to-high teens.

Direct AI Services to Surpass $500B by 2029

Advisory and implementation services present the greatest opportunities

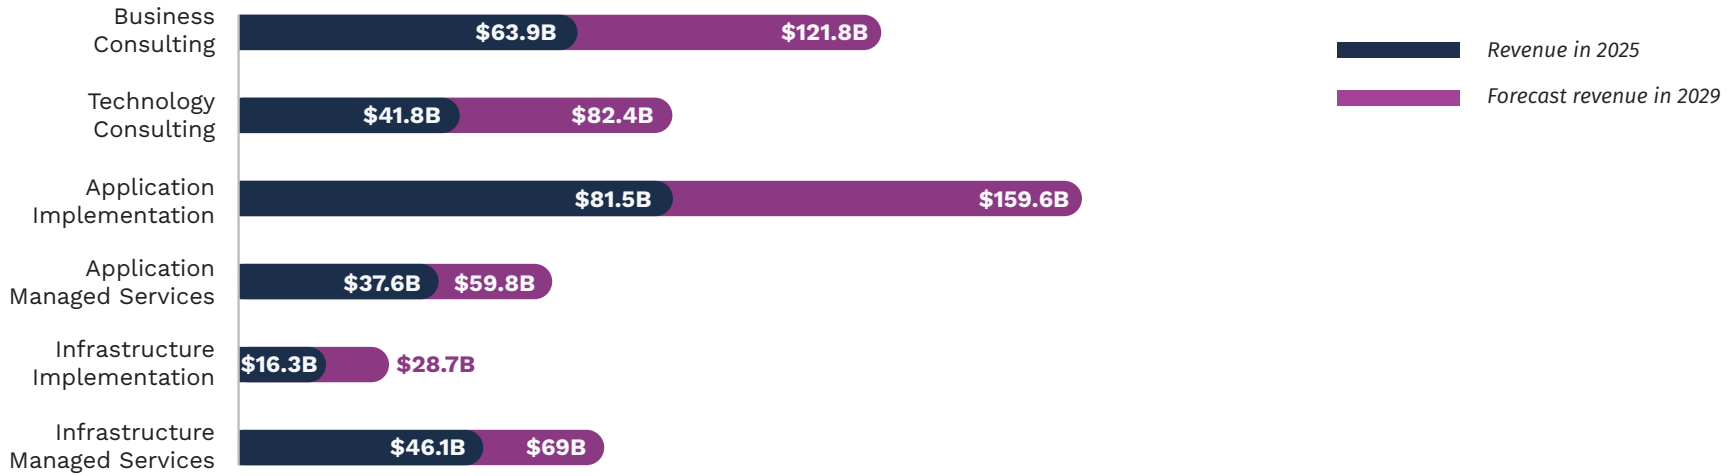

Source: Gartner

In short, while AI technology may be getting all the headlines, making it work within a business still requires trusted partners who can navigate enterprise complexity, redesign processes, reskill teams, and build, govern, and secure AI systems at scale. It's one of the reasons the biggest names in AI are pouring billions into AI-forward partners.

THE CHALLENGE

Yet alongside that opportunity comes a number of challenges for services firms. This includes margin compression, accelerated technology cycles, shifting buyer expectations, and increasing vendor and market volatility. We cover many of these challenges in more detail in our Services Playbook for the AI Era.

Services leaders are well aware they cannot afford to sit this out. They know they must quickly adapt on two fronts: how they operate internally, and where and how they play in the market.

However, they're stuck and overwhelmed, going round and round on questions like:

- Where should I place my bets to capitalize on AI?
- Where should I stop investing to make real progress?
- How do I evolve without tanking the current business?
- Which partners should I double down on and which will be less relevant?
- Should I continue to focus on implementations or a specific channel partner?
- How do I differentiate when everyone is an AI services firm?

FINDING THE RIGHT STRATEGY

Unfortunately, there isn't a single right answer for how to capitalize on this opportunity.

Ideas and advice are being offered to services leaders and founders every day from employees, customers, partners, and investors. But the truth is, we're all in new territory, so identifying and capitalizing on the right strategy takes work and constant adaptation.

It's why we developed this guide. As a framework for approaching the questions listed at left without descending into analysis paralysis. There is a window to capitalize on this opportunity, and if you wait to make the necessary moves, it could be too late.

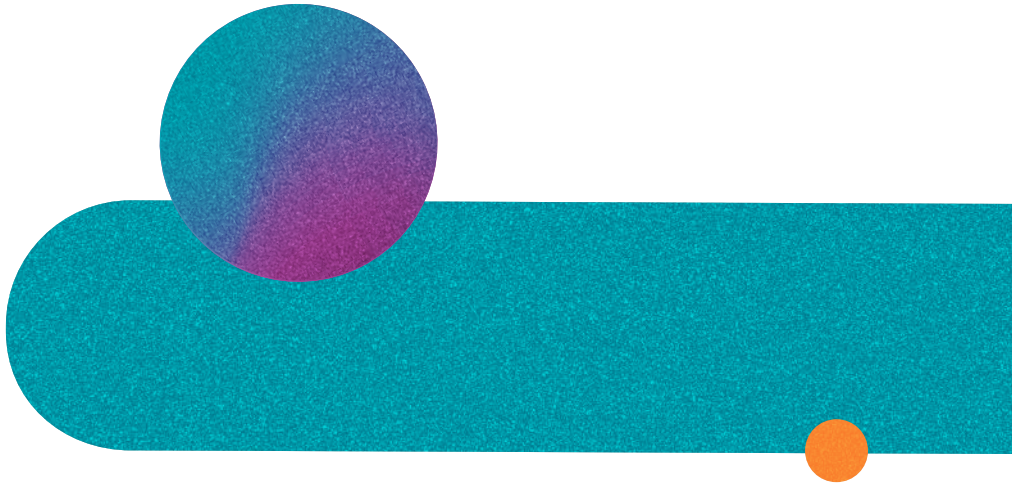

02 / A Basic AI Taxonomy for Services Leaders

AI is often described as a single wave or one monolithic opportunity.

In reality, it spans multiple technologies, approaches, and deployment models, each operating at different tiers of the enterprise IT stack.

Service firms increasingly need to be fluent in AI fundamentals across this entire continuum. Fluency doesn't require technical mastery, but when thinking through strategy, it's important to start with the basics.

The categories of AI

At a high level, AI can be grouped into three broad categories:

Predictive AI: Forecasting and Predicting
Data and algorithms are used to recognize patterns and forecast outcomes. This is largely deterministic, meaning the outcome is entirely determined by the initial inputs and conditions. It includes areas such as credit risk scoring, fraud detection, demand pricing, next-best-offer, and other use cases that have been evolving for years.

Generative AI: Creating and Assisting
Large language models (LLMs), small language models (SLMs) and other multimodal GenAI models are used to generate new content, retrieve knowledge, and act as co-pilots across different roles. This type of AI is non-deterministic, meaning systems and algorithms can produce different outputs with the same input.

Autonomous AI: Acting and Deciding
Also known as agentic AI, this approach relies on agents and teams of agents to act independently; coordinating workflows, triggering transactions, and executing tasks end-to-end with little human intervention. This type of AI is also non-deterministic.

The basic layers of AI

AI capabilities cut across every layer of the enterprise IT stack — from the infrastructure that powers model training, to the data platforms that fuel AI systems, to the applications and workflows that ultimately deliver business value.

It's easy to over-index on visible applications because that's where innovation feels immediate and tangible. However, the performance of these front-end experiences depends greatly on the layers beneath, with each building upon the next. Underinvestment in the foundations or weaknesses anywhere in the stack, such as fragmented data, poor integration or unclear oversight, can cause AI initiatives to fall far short of expectations.

Here is a simple overview of the different layers and how they interrelate.

LAYER	INCLUDES	PURPOSE
Application / Agent Surfaces	Copilots, chat interfaces, embedded AI features in business applications	Delivers AI capabilities directly into business workflows
Orchestration / Control Plane	Observability, workflow engines, API management, agent frameworks, prompt routing	Coordinates model calls, agents and integrates with enterprise systems
Knowledge / Semantic	Vector databases, RAG systems, knowledge graphs	Provides contextual grounding using enterprise knowledge and standardized definitions
Data / Systems of Record	Data ingestion, ETL pipelines, data lakes/ warehouses, enterprise systems of record	Ensures high quality, accessible, governed data to fuel models
Model	Foundational models, custom ML models, finetuned LLMs (industry, subprocess, etc.), embedded models	Delivers predictive or generative intelligence
Infrastructure / Compute	Cloud platforms, GPUs/TPUs, networking, scalable compute	Provides processing power for data storage, training and interface at scale

GOVERNANCE
SECURITY
COMPLIANCE

- Observability
- Monitoring
- Governance
- Compliance
- Identity

Ensures safe, compliant and reliable AI deployment

There is significant overlap between these layers, and as AI evolves, some of them will likely merge. For example, observability (the ability to maintain continuous visibility into how AI systems behave and whether that behavior is acceptable) could be considered part of both the orchestration and governance layers. Similarly, it's also hard to fully detach the knowledge layer from orchestration.

The terminology is also nuanced. Take orchestration, one of the hottest terms within AI these days. RAG orchestration might focus on retrieval, ranking, grounding, and answer generation. Agent orchestration might focus on planning, tool selection, delegation, and memory. And workflow orchestration might emphasize deterministic control of business processes.

We define these layers not to overwhelm you, but to explain that there are multiple facets to any AI strategy or solution. Services companies may choose to specialize in a few of these areas, developing IP and partnerships to stand out in a very crowded space. However, to deploy enterprise-grade solutions, they'll need to develop an understanding of capabilities across most of the layers.

Building an AI strategy starts with understanding the nuances of AI and how AI capabilities span the enterprise IT stack.

03 / The Evolving Vendor Landscape

AI and the vendor landscape around AI are evolving faster than most can keep up.

Given this pace of change and the rampant AI washing happening across vendors, where everyone has repositioned themselves as an AI leader, it can be difficult for anyone to fully understand who all the players are. Much less predict who the winners will be.

THE POWER PLAYERS

This said, a clear set of power players is emerging in the lower layers of the AI stack, which provide the essential ingredients for AI to work: the infrastructure, models, and data.

There is a growing number of **foundational model providers** at the moment, the enterprise market appears to be coalescing around Anthropic, Google, Microsoft and OpenAI. That said, who is at the top of that leaderboard can shift quickly. In early 2025, OpenAI would've been considered the one to beat. A year later, Anthropic is surging in the enterprise.

Then there are the **infrastructure and compute providers**, where chip manufacturers like Nvidia and the hyperscalers — AWS, Google and Microsoft — dominate. These players serve as the power plants and landlords of the high-compute environments that power the AI economy.

Finally, there are **data platforms** that provide the raw materials needed to train the models. The largest data players — Databricks, Oracle and Snowflake — continue to scale, while newer AI-native data players like MotherDuck, Starburst, and Dremio (acquired by SAP in May 2026) are finding their way into specific domains.

Revenue Duel
Anthropic has rocketed ahead of OpenAI in annualized revenue

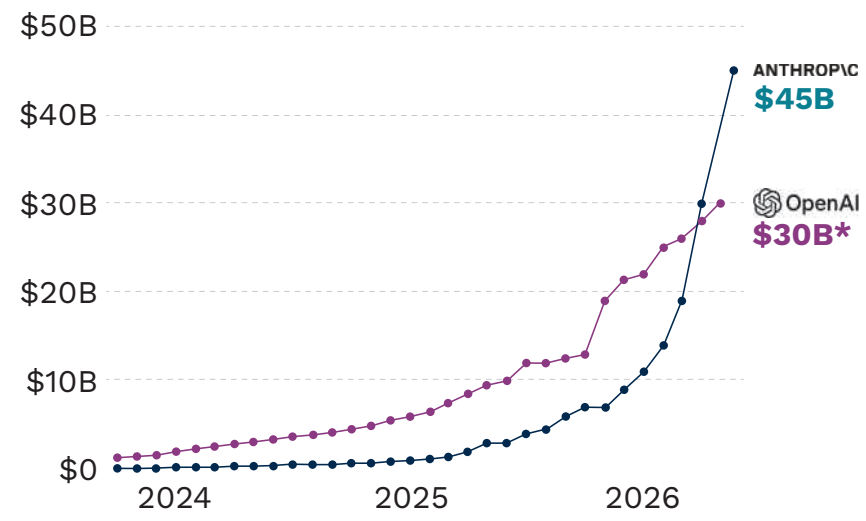

Source: Estimates from The Information, May 2026

APPLICATIONS AND SYSTEMS OF RECORD

At the upper end of the AI stack are the application providers that deliver AI capabilities into business workflows. These include AI-native providers as well as traditional software and SaaS vendors that are building AI capabilities into their platforms.

There is a narrative that AI will render traditional software obsolete. Some software categories are certainly at greater risk of disruption and could be more easily replicated with AI. For example, standalone applications with lower complexity, or those that serve as intermediary interfaces or aggregators.

However, we believe it will be a long time before enterprises have the desire or capability to replace complex business systems with custom AI workflows that internal teams (even with the help of agents) would need to design, build, maintain, extend, and keep compliant.

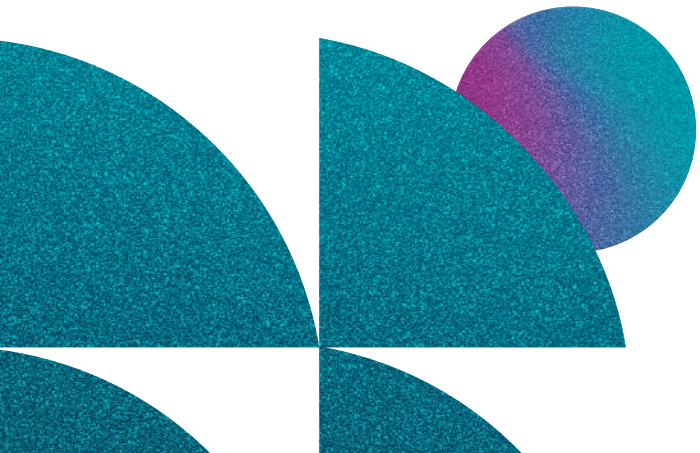

We also believe there is a difference between simple departmental or functional applications and enterprise-grade systems of record ("SORs"). SoRs are mission-critical systems, deeply embedded in operational and often regulated workflows. They encode a business' institutional memory and its data model, and are integrated with sometimes hundreds of downstream systems that depend on this data to generate or track revenue. The more scale, data, integrations, and performance required, the harder these systems will be to replace.

Large enterprises have made massive systems investments in platforms such as Atlassian, Salesforce, SAP, ServiceNow, and Workday over the last few decades. Customers might be cautious about major new investments in these systems in the near term, but they'll want to extend the value of their existing investments as long as possible.

Obviously, these systems must evolve to stay relevant, which is not lost on the vendors. We're seeing greater investment in headless architectures that would expand the lifespan and value of existing systems. SORs are building AI features into platforms to leverage their captive data. And most continue investing into new markets through a mix of build, buy, and partner strategies.

THE ENTERPRISE ENABLERS

Then there are the AI-native vendors building out the middle, enabling layers of the AI stack. These vendors provide contextual grounding of data, and coordinate model calls and orchestrate agents that handle monitoring, permissioning and compliance. They are the vendors that make agentic and autonomous AI in the enterprise possible.

Names like:

- Credo AI, Arize, Fiddler AI, and Vanta for governance and security
- LangChain, OpenClaw, ARIA, and CrewAI for orchestration
- Pinecone and MongoDB for vector search
- Neo4j and Glean for knowledge graphing

The vendors in this middle layer are multiplying almost daily, and it's becoming one of the most the fragmented areas of the enterprise AI stack. We believe many of these tools will consolidate, become acquired or eventually become embedded within broader AI platforms.

However, for now, this is one of the most exciting areas to watch in enterprise AI because it's where AI prototypes progress into real enterprise systems.

RISE OF AGENTIC DOMAIN SPECIALISTS

We're also seeing an explosion of AI-enabled platforms looking to automate specific industry or domain-specific workflows.

These include solutions to automate coding and software engineering (Cursor, Claude, Codex, Harness), systems implementations (Auctor), accounts payable (Vic.ai), customer support (Sierra), and recruiting (Paradox). Just to name a few.

This category also includes industry-specific solutions. A few examples of these vendors include Harvey for legal contract development, Abridge for clinical healthcare documentation, Buildots for construction project management, and Tractable for insurance claims assessment.

Over time, these "services-as-software" platforms have the potential to replace existing services and software.

It's why we are seeing a number of domain-focused consultancies both partnering with these AI-native players to help them speed delivery and differentiate, while some are choosing to develop their own IP here.

A VENDOR LANDSCAPE IN FLUX

The fragmentation and accelerated pace of the AI vendor ecosystem can make it challenging for services firms to know which vendor to build relationships and capabilities around. Especially since many of these newer AI players have relatively nascent partner programs and may be ripe for acquisition as layers merge and the market consolidates.

However, this also represents a significant opportunity for services firms. Those who can guide their clients through this complexity, advise on vendor selection, and provide confidence around interoperability, security, and long-term stability stand to benefit from the confusion.

It may also present an opportunity for firms to be early movers with the right vendors, collaborating on solution development and providing customer feedback to help mature solutions and models for enterprise-grade deployment.

Just know there is a different playbook for partnering with the AI natives. More on that here.

04 / A Framework for Prioritizing AI Service Opportunities

In the cloud era, an entrepreneur might look to build and scale a tech services company off the back of a single software ecosystem. In the AI era, that is becoming increasingly difficult.

One of the biggest challenges is that the vendors themselves are in a state of flux. It's hard to know exactly which platforms will make the leap into the AI era, and which will be relegated to legacy technology that AI builds around. Vendors are also embedding AI directly into their platforms, automating parts of the implementation and support work that once fueled partner margins.

In general, the gravity is shifting from implementation to orchestration work because AI demand does not fit neatly within a single application. Enterprise AI work typically cuts across clouds, data stores, models, governance layers and workflows, which means value increasingly lies in connecting models and proprietary data, process redesign, security, compliance, and change management. That favors firms that can solve business problems end-to-end, not those that only know one vendor's platform.

Having unique expertise and deep relationships with vendors is still critically important. It's just not the only factor.

So how does one determine the right go-forward strategy? How do firms know which AI opportunities are the right bets to pursue, which aren't, and what the right level of investment should be? Especially if they have an existing business to support.

The Tercera Advisory team has been spending time with many tech services firms to help them answer many of the questions posed in this report. While every firm is unique, we find the best opportunities sit at the intersection of:

1. Market demand
2. Customer needs
3. Domain expertise
4. Technology advancements

Strategic Opportunity Framework
From Insight to Impact

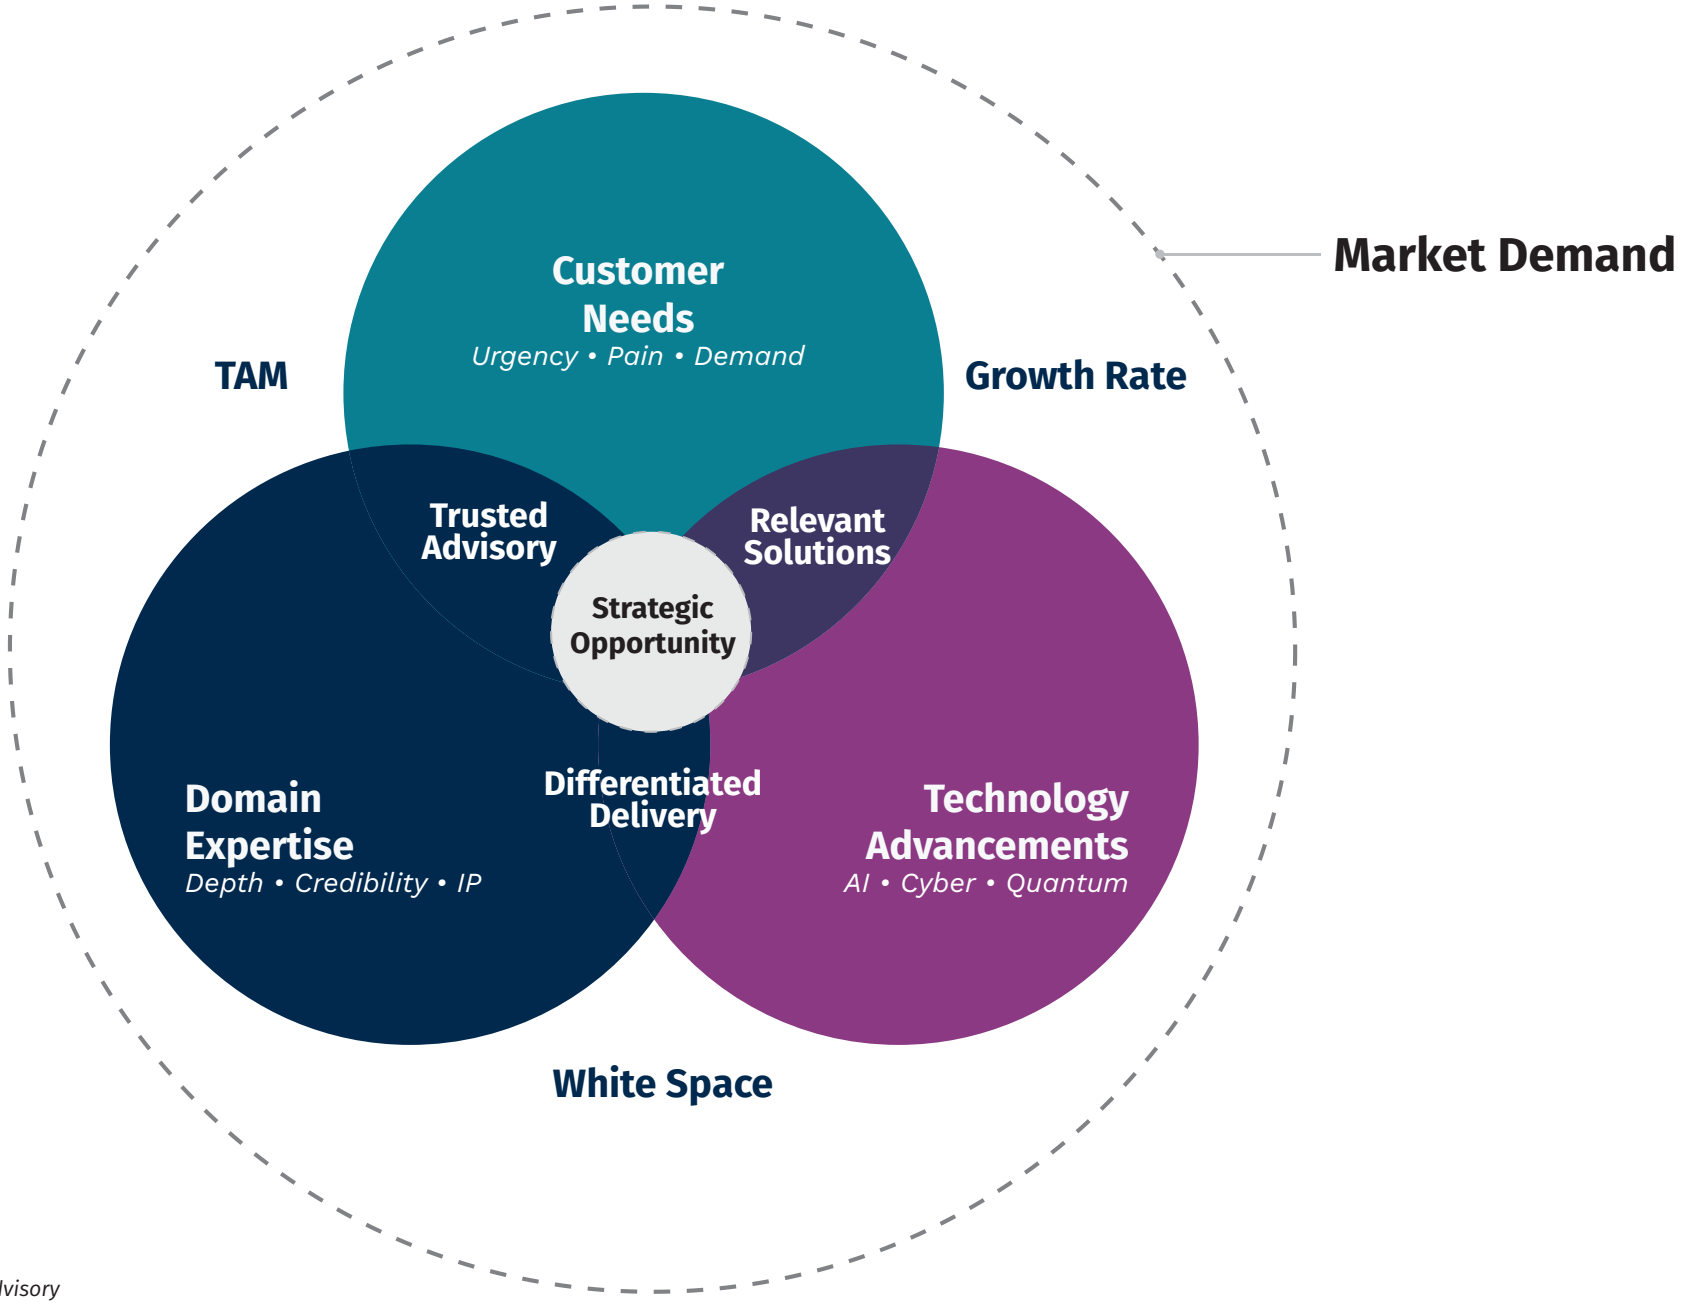

Source: Tercera Advisory

MARKET DEMAND

When evaluating future opportunities, you must first have an understanding of where the market is heading. With AI's research capabilities, it's easier than ever to gather insights to find market sizing data and growth rates, and to understand competitive dynamics.

The tougher part is translating that market intelligence into a differentiated thesis — one that identifies where customer pain is most acute both now and in the future, where incumbent solutions are falling short, and where a services company can build a durable advantage rather than simply chasing the latest wave of demand.

It's also important to understand how AI services will compress traditional services spend. As mentioned in the previous section, software vendors are already automating parts of net-new implementations. Agents are beginning to take on manual code development and testing, while certain aspects of back-office execution across marketing, sales ops, legal, and finance are becoming automated.

However, as these areas shrink, new opportunities will emerge. This happens in every tech cycle.

The important thing is to think about how to balance exposure to areas where AI will compress demand with areas where AI will create new demand. And to build a services model that can evolve as the profit pools shift.

CUSTOMER NEEDS

Great strategies require a deep understanding of customer problems — both seen and unforeseen. You can start with a thesis, but without real customer conversations it's just theory.

Start by conducting bottom-up conversations directly with customers or prospects, as well as with delivery teams and forward-deployed engineers who are close to customer problems and the tech platforms themselves. If you are a channel-driven business, having discovery conversations with your strategic vendor partners is also important, especially if you're dependent on them for leads.

If you're building a firm from scratch, there is freedom to explore a wider universe of customer needs. However, it will take real work to reach real buyers (not just friends in your network) and to understand second- and third-order needs. If you're an existing firm with specific domain expertise and established customer relationships, you have a leg up. You have a built-in focus group.

This kind of customer access is why we believe mid-size and specialized IT services firms have a huge opportunity to take share in the AI era. If they play it right.

The GSIs are investing millions in the pivot to AI, but they are still weighed down by their size and a business model that depends on multi-million-dollar, multi-year programs. Specialist and mid-size firms can move faster, and their very existence depends on speed, which lends a sense of urgency. Those with domain expertise also have a leg up, with a built-in customer base in which trust and deep knowledge of business processes and requirements are already established.

DOMAIN EXPERTISE

Domain expertise has always been important in services, but it becomes even more important in the AI era. A generic AI solution may look impressive in a demo, but fail when it meets messy data, exception-heavy workflows, legacy systems, regulatory obligations, or frontline users who do not operate the way a generic agentic workflow assumes.

Domain experts understand the edge cases. They know which decisions require precision, which processes can tolerate approximation, where automation is risky, and where human oversight is essential.

A team with deep industry knowledge can identify the use cases that matter most, redesign the surrounding process, anticipate stakeholder resistance, and define the right success metrics.

Domain expertise is also one of the clearest ways for services firms to build durable differentiation. As the underlying AI tools become more accessible, technical implementation alone will be harder to defend. The firms that stand out will be those that combine AI fluency with specialized knowledge of a sector, function, or business process.

If you're a firm starting from scratch, it may be tempting to target the domains with larger TAMs and enterprise spend. Industries such as financial services, healthcare, and manufacturing. It's not a bad strategy, but sometimes there is more white space in less sexy industries, or in areas where your founding team has direct expertise.

If you're an existing business, you may be more restricted to domains where you have permission to play. Yet, existing clients could give you a built-in advantage.

Growth adjacencies often sit at the intersection of established domain expertise and emerging needs and technology advancements.

Locking in on your established domain expertise can be surprisingly difficult, especially for horizontal or diversified services firms that don't specialize in a particular type of customer, buyer person, industry, or platform.

A few questions that might help:

- Where do clients already trust us to solve high-stakes problems?
- Which industries, functions, or buyer types show up in our best work?
- What processes, workflows, or systems do we understand better than most?
- Where have we built repeatable assets or IP to solve these problems?
- Where do we have privileged insight or proprietary data that is valuable to customers and would make it hard for a competitor to easily copy?
- Where can our current expertise create an unfair advantage if we apply AI?

TECHNOLOGY ADVANCEMENTS

This aspect of a strategy may be the most difficult to hone in on, especially given how quickly the underlying technology advances. It is hard enough to understand what AI can do today, let alone what it may be capable of in a few months or a few years. Models are improving across reasoning, multimodal understanding, code generation, memory, and agentic execution. All of these can materially change what work can be automated, augmented, or rebuilt from the ground up.

Part of the challenge is that advancements do not move in a straight line. A capability that feels unreliable today may become production-ready in a few months once models improve, context windows expand, inference costs fall, or better orchestration and governance layers emerge.

At the same time, breakthroughs in adjacent areas — cybersecurity, edge computing, robotics, synthetic data, quantum computing, and AI infrastructure — can further change the art of the possible. For example, AI-native security could reshape how enterprises monitor and respond to threats; robotics could extend AI-driven automation into the physical world; and quantum computing could eventually alter aspects of cryptography, optimization, and scientific discovery.

Given this uncertainty, services firms need to build adaptability into the strategy itself. That means creating mechanisms to continually scan the technical and vendor landscape, and to test emerging capabilities. It's helpful to have technologists and AI-native practitioners that are close to the frontier involved in strategy discussions, and to build internal or offshore innovation labs where teams can experiment with new tools before customers ask for them.

The goal is not to predict every breakthrough, but to develop the organizational muscle to recognize meaningful shifts early, separate durable opportunities from hype, and reposition the business before customer demand and services profit pools move.

Now is the critical window to take action.

AI is both the wave and the undertow for services firms. It will create massive new demand, while pulling away parts of the business that once felt predictable and protected. That is the paradox.

The firms that win in this environment aren't those that chase every shiny AI opportunity, nor the ones that wait for the dust to settle. They are those that choose to keep pace with how the technology and vendors are evolving. Those that move with urgency, but are deliberate about the opportunities they pursue.

TERCERA

Tercera is an investment and advisory firm that helps AI-native and traditional tech services firms grow, evolve, and win in the AI era. As experienced services operators and investors, we understand what it takes to build exceptionally valuable services businesses, and we bring the capital, counsel, and connections to turn vision into reality.

If you are looking for external help defining and making more meaningful progress on your AI strategy, we'd love to hear from you.

Contact us

Subscribe to our newsletter

Learn more about our Advisory services