SVS THOUGHT LEADERSHIP · AI TRANSFORMATION STUDIO

Lead the AI-Business Transformation

From AI Adoption to a New Way of Running the Business — with full research references.

From AI Adoption to a New Way of Running the Business

The current AI debate is still dominated by tools. Which model is best? Which copilot should we deploy? How many employees are using AI? How much time can we save?

These are relevant questions, but they miss the larger transformation.

AI is changing the economics of what companies can build and how they can operate. The provocative assumption of producing twenty times more software output with the same budget illustrates the magnitude of that shift. Capabilities that were previously too expensive to develop or maintain can suddenly become economically viable. Digital services can enrich physical products. Traditional services can become continuous and data-driven. Planning can move from periodic cycles toward real-time optimization. Decisions that previously required several organizational handovers can increasingly be synchronized across functions and systems.

The first-order effect is productivity. The much more important second-order effect is **business transformation**.

This distinction matters because access to AI itself will not create a lasting competitive advantage. The underlying models, cloud infrastructure, development tools and commercial applications are becoming broadly accessible. If every competitor can buy similar technology, advantage has to come from somewhere else.

Strategy has dealt with this question long before AI. Resource-based and capability-based strategy argue that sustainable advantage comes from resources, knowledge, routines and organizational capabilities that competitors cannot easily reproduce. AI does not invalidate that logic. It makes it far more consequential.

The emerging competitive race is therefore not about who adopts the most AI. It is about who can turn AI into **distinctive business capabilities** faster, embed those capabilities into the operating model, and continuously improve them.

Two external bodies of work help clarify this transformation. McKinsey gives us important evidence about the **WHAT**: where AI creates value, which business domains matter most and what distinguishes leading AI transformations. Palantir gives us useful evidence about the **HOW**: how data, business logic, workflows, agents and governance can be connected into an operational system that turns decisions into action.

But neither McKinsey nor Palantir is the thesis of this paper.

The thought leadership here is SVS's synthesis: **Lead the Transformation → Build the Capability → Run the AI-Operations**. This is our framework for connecting strategic ambition, capability building and AI-enabled execution. McKinsey helps establish why focused business transformation matters. Palantir helps illustrate how an operational architecture can make that transformation executable. SVS defines the management logic that connects the two.

McKinsey research on leading AI transformations supports this distinction. The strongest performers do not spread hundreds of disconnected AI experiments throughout the company. They concentrate on a small number of economically important business domains, transform them deeply and connect technology investments to measurable business outcomes. In McKinsey's research on AI leaders, transformations produced an average 20 percent EBITDA improvement and typically focused on only one to three business domains. [1]

The transformation challenge therefore has three connected stages:

Lead the Transformation → Build the Capability → Run the AI-Operations

These stages describe more than an implementation methodology. Together, they represent the SVS perspective on the journey from AI as a productivity technology to AI as part of the operating fabric of the enterprise.

---

1. Lead the Transformation

AI creates an economic shift before it creates a technology shift

The starting point cannot be technology. Leadership first has to understand how AI changes the economics of its industry, its customers and its own business model. When software, analysis, coordination and increasingly decision execution become dramatically cheaper and faster, assumptions that shaped yesterday's business model begin to break down.

A manufacturer may discover that selling equipment is no longer the most attractive model when connected equipment, predictive services and lifecycle optimization become feasible. A service provider may move from selling expert hours toward continuously delivered digital capabilities. A distributor may discover that pricing, inventory and customer demand can be synchronized rather than periodically reconciled. A complex B2B provider may be able to turn individual sales excellence into an AI-augmented organizational capability that guides every seller through a consistent decision process.

The leadership task is therefore not to ask where AI can be inserted into the existing company. It is to understand **what company should exist when the old technological constraints no longer apply**.

This is why broad AI adoption is such a weak strategic starting point. Individual productivity gains are valuable, but they are diffuse and increasingly easy to replicate. Competitive advantage begins

when the company identifies the few economic leverage points where a significant improvement changes its competitive position.

McKinsey's research into successful AI transformations strongly supports this focus. Its analysis emphasizes that technology itself is broadly available; differentiation comes from the capabilities that allow companies to apply it to important business problems faster and more effectively than competitors. The same research argues that business leaders, rather than IT functions alone, have to own this reimagination because the issue is business performance, not technology deployment. [2]

The SVS contribution is to turn this evidence into a transformation sequence. The question is not only where AI creates value. It is how leadership converts that opportunity into a deliberate choice about the capabilities the company must build and the operating system it must eventually run.

Transformation leadership means choosing where to become distinctive

This leads to a different kind of transformation roadmap. The roadmap is not a collection of AI use cases. It is a deliberate choice about **where the enterprise intends to become disproportionately good**.

Every business has domains where performance matters more than elsewhere. Yield and throughput may dominate the economics of a mining operation. Supply-chain integration may determine working capital and availability in another business. Pricing may create disproportionate margin leverage in retail. Equipment uptime may define customer value in an industrial service model.

The same principle applies to less asset-intensive businesses. Customer acquisition, complex solution selling, innovation, project delivery or advisory services can all become strategically important domains.

The transformation roadmap starts by connecting external market forces with internal capabilities and competitive ambition. It defines the future strategic position of the company, identifies the domains that must support that position and determines what new capabilities are required. Only then does technology enter the equation.

This reverses the usual AI discussion. Instead of starting with what a model can do, leadership starts with what the company needs to become. This is the first distinctive SVS proposition: AI transformation should be led as a strategic redesign of the business, not as a technology adoption program.

De-risk the transformation before scaling it

There is another important implication. A bold strategic ambition should not become a multiyear leap of faith. Transformation needs a mechanism for converting conviction into evidence.

Palantir's AIP Bootcamp provides an instructive example. Rather than beginning every engagement with a large-scale technology program, Palantir brings business users and engineers together around actual operational problems and actual company data. Palantir's 2023 business update

described the original Bootcamp model as moving from zero to a working use case in five days or less. Its 2025 Annual Report continues to describe Bootcamps as a mechanism for delivering real workflows on customer data within days. [3][4]

The significance is larger than rapid prototyping. A concentrated transformation sprint tests several assumptions at once. Is the problem economically important? Can the envisioned capability improve it? Is the necessary data available? Can the technology work? Will users accept the new workflow? What organizational obstacles become visible?

That makes the bootcamp or transformation studio a **de-risking mechanism** between strategy and implementation.

Transformation leadership therefore combines ambition with disciplined validation. The company defines where it wants to create strategic distance, but it earns the right to scale by proving the underlying capability.

In the SVS framework, this is the bridge from the WHAT to the HOW: a focused, evidence-based way to test whether a strategic ambition can become a working business capability.

---

2. Build the Capability

The unit of transformation is the capability, not the AI use case

Once strategic direction is clear, the next challenge is turning ambition into something the organization can actually do.

This is the capability shift.

A business capability is more than a process and more than a software application. It combines people, expertise, workflows, decision rights, data, business logic, technology, governance and management routines to reliably produce a business outcome. This is why isolated AI use cases so often disappoint.

Consider predictive maintenance. Predicting a machine failure is technically useful but economically incomplete. Value appears only when the surrounding capability can respond. Maintenance priorities need to change. Technician availability must be understood. Spare parts have to be located or ordered. Production implications need to be considered. Customer commitments may have to be adjusted. The intervention must be executed and its result captured.

Improving the prediction without redesigning the surrounding system creates intelligence without execution. The capability approach changes the perspective. Instead of asking how AI improves one task, the organization redesigns how the complete business outcome is produced.

This is also why the strongest AI transformations focus on end-to-end business domains. McKinsey's research on agentic AI distinguishes individual augmentation and task automation from agentic workflows and systems that coordinate multiple activities. The latter can influence

throughput, cycle time, margin, working capital and resource utilization precisely because they optimize across the workflow rather than within one isolated task. [5]

McKinsey helps establish the WHAT: value is created when AI is applied to economically important domains and connected to measurable outcomes. The SVS framework adds the management implication: those domains must be treated as capabilities to be built, not as disconnected use cases to be deployed.

The capability shift requires an operating shift

Building a new capability inevitably exposes weaknesses in the existing operating model.

Traditional organizations are designed around functions. Sales owns sales. Operations owns operations. IT owns technology. Finance owns financial control. Data is frequently distributed across several of them. Risk and compliance often enter the process through separate approval mechanisms.

This structure can work when information moves slowly and decisions can be decomposed into functional tasks. It becomes a serious disadvantage when the competitive opportunity comes from synchronizing decisions in real time.

Complex operations contain enormous amounts of stranded value between functions. Demand changes while production plans remain fixed. Sales commitments are made without current operational constraints. Service technicians discover opportunities that do not reach commercial teams. Inventory decisions optimize availability without understanding customer value. Insights sit in dashboards while operational systems continue doing what they did before.

AI makes these gaps more visible because the technological cycle becomes faster while the organizational cycle remains slow. The operating model therefore has to evolve together with the capability.

Business, technology, data and relevant functional expertise have to come together around the outcome. Decision rights need to support fast action. The people responsible for the domain require end-to-end accountability. Technology platforms have to enable reuse rather than create new dependencies. Performance measures need to reflect business outcomes rather than completion of projects.

The objective is not simply agility. It is to reduce the latency between **signal, decision and action**.

McKinsey's research describes this as an operating-model issue because technology is increasingly available to everyone; the differentiator becomes how quickly and effectively the organization can use it. Cross-functional teams, business ownership, reusable technology platforms, consumable data and reconfigured end-to-end processes reinforce one another. [2]

The SVS perspective makes this operating shift explicit. A capability is not complete when a model works. It is complete when the organization can repeatedly use the model, act on its output, govern the resulting decisions and improve the outcome over time.

Distinctive capabilities are where the moat begins

The capability perspective also clarifies what companies should build themselves.

There is little strategic reason to recreate commodity infrastructure. Standard finance software, generic productivity applications or basic cloud services normally create more value when purchased.

But a system that expresses what the company uniquely knows and does may deserve very different treatment.

A proprietary pricing capability can incorporate years of company-specific transaction data and market knowledge. An equipment-performance capability can encode engineering expertise unavailable to competitors. A customer-development system can embed a company's strategic sales methodology, accumulated account intelligence and decision logic. The technology becomes valuable because it carries the capability.

This is where the resource-based logic of strategy connects with AI. The defensible asset is not the foundation model. It is the combination of **domain knowledge, proprietary data, workflows, business logic, learning and execution** that surrounds it. And unlike a standalone application, such a capability compounds.

Every decision generates additional information. Every execution produces an outcome. Every outcome improves the understanding of what works. Reusable objects, rules, workflows and models can support the next application. Over time, the company is no longer merely implementing AI solutions. It is building an increasingly rich digital representation of how it competes. That prepares the ground for the third shift.

---

3. Run the AI-Operations

From AI-enabled workflows to an executable operating system

The third stage begins when distinctive capabilities become digitally explicit enough that humans and AI agents can operate them together.

This is the shift from **building AI-enabled capabilities to running AI-Operations**.

The term matters. Running AI-Operations is not the same as running an AI platform or maintaining machine-learning models. It means that important parts of the business increasingly operate through AI-augmented, closed-loop systems.

Palantir provides a clear enterprise-scale example of this architectural direction. Palantir's platform documentation describes AIP, Foundry and Apollo as components of an **Enterprise Operating System**. At its center is the Ontology, which represents not simply enterprise data but the objects, relationships, logic, permissions and actions through which the organization operates. Palantir

describes the Ontology as an operational representation of the enterprise's interconnected decisions rather than merely a semantic representation of its data. [6]

This distinction is fundamental as traditional enterprise technology is largely organized around applications and databases. One application contains customers, another orders, another assets, another financial information. Humans understand how these fragmented elements fit together and coordinate them through processes. An operational ontology makes this context machine-readable.

A customer, machine, order, contract or opportunity becomes an object with defined relationships to other objects. Business logic can operate on those objects. Actions specify how their state can be changed. Permissions determine who or what is allowed to take an action. Applications and agents can work on the same operational representation.

Palantir's documentation explicitly includes actions as part of the Ontology and describes them as structured mechanisms for changing objects and connecting decisions back into existing systems. AIP then adds agents, AI functions, evaluations and automations on top of this operational environment. [7][8]

The consequence is profound. AI no longer merely knows something about the company. It can participate in how the company operates.

Palantir helps make the HOW tangible. But the SVS proposition is broader than adopting Palantir's architecture. The essential question is whether the enterprise has created the strategic capabilities, operating model and governance required for AI to act meaningfully within the business.

Close the loop from information to action

Traditional digital systems have become extraordinarily good at collecting and displaying information. But the chain frequently stops at insight.

Data becomes a report. The report becomes a dashboard. The dashboard shows a deviation. A human then has to understand the situation, coordinate with several functions, decide what to do and manually initiate the required actions. AI-Operations extend this chain.

Data → Operational Context → State → Analysis → Decision → Action → Outcome → Updated State

An AI agent can detect a change in the operating environment, interpret it using the company's business context, evaluate possible actions and recommend the next step. Depending on the level of autonomy permitted, a human may approve that action or the agent may execute it directly. The resulting change becomes part of the new operational state. The loop starts again.

That is the decisive difference between an AI assistant and an AI-enabled operating system. The first produces output. The second influences **state**.

This principle is already visible in Palantir's architecture. Foundry's Ontology allows applications to capture decisions and write them back into the operational environment. AIP enables agents to work with those objects, functions and actions. Palantir's application-development tools are designed to enrich the same consistent set of data and logical assets so that knowledge can

compound as operational workflows expand across the enterprise. [7][9] This is the technological foundation for closed-loop operations.

In the SVS framework, however, AI-Operations are not defined by the presence of agents alone. They are defined by the existence of a governed, measurable and continuously improving loop between business signals, decisions, actions and outcomes.

AI-Operations do not mean removing humans

An autonomous enterprise should not be confused with an organization in which AI is allowed to make every decision independently. Autonomy is a design variable.

Some actions can be fully automated because they are reversible, low-risk and well bounded. Others require human approval. Some decisions need multiple authorization levels. Sensitive actions may always remain under human control.

The real breakthrough is not removing humans from the loop. It is **designing the loop deliberately**.

Humans should concentrate where judgment, accountability, creativity and exception handling create the most value. Agents can handle monitoring, analysis, coordination, repetitive decision preparation and increasingly execution within clearly defined boundaries. This makes governance part of the operating architecture rather than an external compliance layer.

Permissions, security, observability, evaluations, audit trails, escalation rules and limits on agent autonomy determine the right to act. Palantir's architecture combines AI development with existing security and governance and provides evaluation and observability capabilities for deployed agents. [8] This is critical because enterprise AI cannot scale on enthusiasm alone. Trust has to be engineered.

The SVS view is that governance is not a brake on transformation. It is one of the design disciplines that makes transformation scalable.

From Applications to an AI-Agent Platform Business

The move toward AI-Operations also changes the role of technology inside the company. Historically, most enterprises accumulated applications. Each solved a specific problem and created another repository of data, logic and workflows. Integration became the permanent tax of enterprise IT. An AI-agent platform creates a different possibility.

Core business objects, knowledge, rules, capabilities and actions can become reusable across many operational workflows. Instead of every AI application beginning from scratch, new capabilities build on a shared operational foundation. This creates a compounding effect.

A customer object developed for one workflow can be reused by another. A capability for evaluating risk can support multiple processes. A workflow that captures a human decision can

improve the context available to future agents. A security or governance rule can be applied across applications rather than implemented repeatedly.

Palantir describes this architecture as allowing knowledge to compound as workflows scale through the organization. [9]

This is not merely an IT architecture decision. It is a **platform business shift**. The company gradually moves from operating a collection of applications toward operating a shared digital representation of its distinctive business capabilities. New digital products, internal operating systems and customer-facing services can all build on this foundation.

For some companies, those capabilities may ultimately become products themselves. An industrial manufacturer can expose equipment intelligence as a customer service. A professional-services business can encode parts of its methodology into a digital operating system. A solution provider can turn accumulated domain expertise into an AI-augmented execution platform. AI therefore changes not only how the company works. It can change **what the company sells**.

The SVS implication is that platform thinking should follow capability strategy, not replace it. A platform becomes strategically valuable when it compounds distinctive business knowledge and accelerates the execution of the capabilities that matter most.

One Transformation – Three Connected Shifts

The three stages should not be understood as separate initiatives. They are one transformation.

Lead the Transformation addresses the economic shift and the leadership challenge. It determines where AI changes the competitive environment, which strategic position the company wants to occupy and which business domains must become distinctive.

Build the Capability translates that ambition into a new operating reality. People, processes, data, decision rights, technology and domain knowledge are designed as an integrated system capable of producing superior business outcomes repeatedly.

Run the AI-Operations makes those capabilities increasingly executable. AI agents operate within the company's business context, connect decisions with actions and outcomes, and create closed learning loops under defined governance and control.

This is the central SVS thought leadership: the three shifts are not a sequence of technology deployments. They are a management system for turning AI into competitive capability.

McKinsey provides evidence for the WHAT: focus on economically important domains, connect AI to business outcomes and redesign the organization around value creation.

Palantir provides evidence for the HOW: represent business context operationally, connect agents to governed actions and create reusable systems in which knowledge compounds.

SVS provides the synthesis: **Lead the Transformation, Build the Capability and Run the AI-Operations**. The framework explains how leadership ambition becomes a distinctive capability, how that capability changes the operating model and how the resulting system can increasingly execute and learn through AI.

This also explains why transformation cannot begin with an AI platform. Without strategic direction, the platform automates whatever happens to be there. Without capability design, it digitizes fragmented processes. Without operating-model change, it accelerates individual activities while organizational bottlenecks remain intact. Technology creates leverage only after the organization understands **what it wants to leverage**.

The transformation therefore moves from the outside in and then from the inside out: from market and business strategy to distinctive capability, from capability to operating model, and from operating model into an AI-agent execution architecture.

The Autonomous Enterprise Is an Outcome, Not the Starting Point

There is a temptation to make the Autonomous Enterprise the objective. That puts the technology ahead of the business.

The more useful ambition is a company that continuously improves its ability to create customer value and outperform competitors. Greater operational autonomy becomes a consequence when it improves that objective.

Some capabilities may eventually become highly autonomous. Others will remain predominantly human but become dramatically better through AI-supported intelligence and coordination. The right balance will differ by domain, risk, economics and competitive strategy.

The destination is therefore not a company without people. It is a company in which **people, AI agents, data and technology operate as one increasingly coherent execution system**.

Such a company can sense changes faster, understand their implications earlier, make better decisions, execute them with less organizational friction and learn directly from the results. That is where the productivity story ends and the strategic story begins. AI may make every company more efficient.

McKinsey helps us understand **what** leading AI transformations focus on. Palantir helps us see **how** AI can become operational. SVS provides the thought leadership that connects these insights into a coherent transformation model.

The companies that **Lead the Transformation, Build the Capability and Run the AI-Operations** can use AI not merely to become more efficient, but to become fundamentally better businesses.

Sources

[1] McKinsey & Company, “The State of AI: How Organizations Are Rewiring to Capture Value,” 2025. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

[2] McKinsey & Company, “The AI Transformation: How to Lead Your Organization into the Future,” 2024. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-ai-transformation

[3] Palantir Technologies, “Q4 2023 Business Update,” February 2024. https://investors.palantir.com/files/Palantir%20Q4%202023%20Business%20Update.pdf

[4] Palantir Technologies, “2025 Annual Report,” 2026. https://investors.palantir.com/financials/annual-reports-and-proxies/default.aspx

[5] McKinsey & Company, “The Agentic Organization: Operating Model Transformation for the AI Age,” 2025. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-agentic-organization

[6] Palantir Technologies, “Platform Architecture,” Palantir Documentation. https://www.palantir.com/docs/foundry/architecture-center/platforms

[7] Palantir Technologies, “Ontology Core Concepts,” Palantir Documentation. https://www.palantir.com/docs/foundry/ontology/core-concepts

[8] Palantir Technologies, “AIP Architecture,” Palantir Documentation. https://www.palantir.com/docs/foundry/architecture-center/aip-architecture

[9] Palantir Technologies, “Introductory Concepts,” Palantir Documentation. https://www.palantir.com/docs/foundry/getting-started/introductory-concepts