



**SHI**

**Building your strategic  
AI platform**

# Table of Contents:

<b>Introduction: From scattered experiments to a production-ready AI engine</b> .....	<b>3</b>
<b>Strategic planning: Begin with the end in mind</b> .....	<b>5</b>
The use-case factory – capture ideas, score ROI .....	5
Data: “ready enough” today, better tomorrow .....	6
Rent, buy, or build? .....	7
<b>Organizational readiness: Create an AI council and a learning culture</b> .....	<b>9</b>
<b>Technical platform foundations: Secure, cost-aware, and scalable</b> .....	<b>11</b>
Security first .....	11
Cost management .....	12
Prepare to scale .....	12
<b>Quality and safety: New controls for new capabilities</b> .....	<b>13</b>
Proven playbooks .....	13
Five frontline safeguards .....	14
<b>Implementation approach: From experiment to enterprise scale</b> .....	<b>15</b>
<b>Outcomes: From prototypes to a production-ready AI engine</b> .....	<b>16</b>
<b>About SHI’s AI Readiness Workshop</b> .....	<b>17</b>
<b>About SHI</b> .....	<b>18</b>

# Introduction: From scattered experiments to a production-ready AI engine

Most enterprises have experimented with generative AI pilots, but few have turned these into production services that employees trust, customers notice, and finance can measure. According to Cisco's 2025 AI Readiness Index, eight in ten organizations say the urgency to demonstrate tangible ROI has risen sharply. Yet, only about 13% qualify as "Pacesetters" -- organizations with a disciplined, system-level AI readiness. This highlights a significant gap between ambition and preparedness.<sup>1</sup>

Why the gap? Real advantage doesn't come from sprinkling AI across yesterday's IT stack. It comes from a strategic AI platform: an environment where data, compute, security, networking, governance, and talent converge so AI use cases can be prototyped in days, scaled in weeks, and can be fine-tuned or retired quickly when the next breakthrough arrives.

<sup>1</sup> Cisco. [Cisco AI Readiness Index 2025: Realizing the Value of AI](#). Oct 2025.




Designing that platform requires five essential success factors, which SHI has identified through extensive work with enterprises like yours:

1. **Strategic planning** – What are our goals, policies, and guardrails?
2. **Organizational readiness** – Who owns what, RACI model, and how do we bring the workforce along?
3. **Infrastructure foundation** – What must we rent, buy, or build to explore, prove, and scale?
4. **Quality and safety** – How do we stay legal, ethical, and on-brand when AI helps us create content and make decisions?
5. **Implementation approach** – How do we start small, learn fast, and scale only what works?

In this ebook, we'll explore these five success factors in detail, drawing on SHI's deep expertise in building and scaling AI platforms to help you bridge the gap from experiment to production and achieve measurable results.

<sup>2</sup> Cisco, "[Cisco Study: CEOs Embrace AI, But Knowledge Gaps Threaten Strategic Decisions and Growth](#)," 2025



**“Eventually, there will be only two kinds of companies: Those that are AI companies, and those that are irrelevant.”<sup>2</sup>**

**– Jeetu Patel,**  
Chief Product Officer, Cisco



# Strategic planning: Begin with the end in mind

Generative AI's flexibility creates opportunity for huge wins, but risks a scattered approach. A good strategy clarifies why your organization is betting on AI, which scarce resources you'll invest, and how often you'll revisit priorities as models, regulations, and competition evolve. Write it all down and treat it as a living roadmap – update quarterly and tie it to business outcomes executives already track.

Building your strategic AI platform

## The use-case factory – capture ideas, score ROI

To cast a wide net and focus on solid opportunities, run a lightweight Use-Case Factory:

1. Employees submit a short form: problem, desired outcome, and available data.
2. Score ideas in key areas such as business impact and data readiness.
3. Highest-scoring items enter a prioritized queue; others are parked for future consideration.

**AI Readiness and Use Case Discovery Workshops**

## Data: “Ready enough” today, better tomorrow.

AI initiatives depend on data, but progress shouldn't wait for perfection.

1. **Start with “clean-enough” data:** Prioritize use cases with accessible information, including unstructured sources like PDFs, emails, spreadsheets, and video that modern LLMs can readily consume.
2. **Let early prototypes reveal gaps.** Every prompt stress-tests data quality. Capture issues as insights, without slowing execution.
3. **Fund improvement in parallel.** While quick wins run, identify datasets for next-wave value, invest in schema harmonization, personally identifiable information (PII) cataloging, and vector indexing to enable scale.

SHI® AI & Cyber Labs provides a targeted six-week prototype to help your organization surface data gaps before they impede momentum.

Test and validate your AI use cases and solutions with SHI AI & Cyber Labs

“Most proofs of concept fail because the data is dirty, not because the model is wrong. We told one manufacturer to pause a \$90 million GPU order after a two-week sprint surfaced holes in their order data.”

— Brian O'Connor,  
Solution Architect, SHI



## Rent, buy, or build?

AI adoption is rarely an either/or decision. The most effective strategies use the right mix of rented, purchased, and built capabilities, chosen based on speed, scale, and strategic value.

Start by renting to enable AI quickly within existing tools and build organizational fluency. For differentiated needs, build incrementally using cloud APIs before committing to large vendor purchases, which may be premature in a fast-moving market.

Option	What you get	When it fits
<b>Rent:</b> Easy-to-adopt, incremental capabilities such as SaaS copilots, cloud AI models.	Instant productivity; vendor-managed updates.	Knowledge workers need help today with acceptable security.
<b>Buy:</b> Specialized full-feature AI applications, deploying your own models.	Deep features out-of-box.	Department-wide process with quick ROI expectations.
<b>Build:</b> Custom solutions on your stack.	Control, data sovereignty, lower cost at scale.	Unique data, in-house talent, clear governance.

Building your strategic AI platform

For each component that you rent or buy, fill out a vendor assessment checklist. Here are some good questions to get you started:

- Does the provider publish a shared responsibility matrix (what they secure vs. what you secure)?
- Can you view model cards that explain their AI training data, freshness date, and bias tests?
- Do they provide assurance that they will not use your data to train models used by others?
- Do they provide short or zero data retention and bring-your-own-key encryption?
- Is there a clear migration path from this vendor to whatever comes next?

## Strategic AI investments: Rent, build, or buy?



Building your strategic AI platform

Building a robust AI infrastructure can be complex and time-consuming for your IT teams. It doesn't have to be. At SHI, we have the know-how to simplify your AI deployment process. With over 35 years of experience integrating complex data center and cloud solutions, we offer comprehensive services that span all major AI hardware, software, and hybrid cloud AI architectures. Our pre-configured, pretested, and hybrid cloud architectures provide a low-risk, streamlined approach to integration. You can count on our certified AI engineers, architects and technicians to [streamline your AI deployment](#).

8



# Organizational readiness: Create an AI council and a learning culture

Aligning your organization well is far more important than the technology you use. That's why your first platform component is an AI council empowered to decide where AI fits, how to manage risk, and when to abandon unsuccessful attempts.



Building your strategic AI platform

To be most effective, the council should include all these roles and responsibilities:

Role	Core responsibility
Executive sponsor	Budget authority, issue resolution, KPI enforcement
AI program lead	Council leadership, roadmap ownership, outcome reporting
Product owners	Use case prioritization, business value validation
Data/ML engineers	Data governance, model deployment, and monitoring
Security team	Threat modeling, policy enforcement, compliance
Domain champions	Front-line adoption, feedback collection, pattern sharing

Your AI council should meet monthly to evaluate experiments, kill low-ROI initiatives, and allocate resources for the next cycle.

Creating an AI learning culture is equally vital because, regardless of industry, AI literacy is the strongest predictor of successful adoption. Start with a secure “ChatGPT-like” playground with prompt logging and cost controls. Measure progress quarterly, using the Meta AI Literacy Scale (MAILS),<sup>5</sup> to assess employee competencies and growth.

Create an “AI Town Center” (in Teams or SharePoint) with quick-start guides, prompt templates, and support channels. Enlist champions to maintain these resources and escalate risks early. When one SHI client business unit installed trained champions, weekly active AI use jumped 60 percent over the course of six months, while the average cost per user fell 35 percent. This shows the real-world impact of a structured approach to AI literacy and adoption, which SHI helps organizations put into practice.

<sup>5</sup>Wienrich, C. et al., [MAILS - Meta AI Literacy Scale: Development and Testing of an AI Literacy Questionnaire Based on Well-Founded Competency Models and Psychological Change- and Meta-Competencies](#), ArXiv Paper 2302.09319, February 2023.



# Technical platform foundations: Secure, cost-aware, and scalable

Your technical platform must prioritize security, cost management, and scalability — in that order.



## Security first

AI tools create new data flows that can expose sensitive information. Employees need clear guidelines to avoid inadvertently exposing that information to tools that won't properly secure it. "Most employees are not malicious — they just need a guide to doing the right thing, with guardrails to prevent missteps," notes Aaron Richmond, Senior Solutions architect at SHI. Cloud model vendors might use your data to train their models, potentially exposing proprietary information later.

To prevent this, implement:


- Clear data-handling policies with technical enforcement.
- Personally Identifiable Information (PII) detection and redaction before the data leaves your control.
- Regular security testing of AI systems and integrations.
- Careful evaluation of vendor shared-responsibility frameworks.

## Cost management

As with cloud computing infrastructure, AI costs can catch organizations off guard if not carefully managed. “Understanding AI spending is critical for success,” explains Shawn Cox, AI Product Manager at SHI. “Organizations need visibility into both value created and costs incurred.”

Unlike predictable IT costs, AI expenses vary with usage patterns. Every conversation incurs token costs that accumulate quickly. It isn’t uncommon for each turn of an AI conversation to cost \$0.10 to \$0.25. Multiply that by thousands of employees making dozens of turns each day, and it’s real money. One retailer reduced its AI cloud expenditure by 40 percent simply by routing easier queries to smaller models. This example highlights the significant savings achievable with smart AI resource management.

Start with a focus on visibility rather than optimization, which can come later. Effective management requires usage monitoring with specific attribution to applications and people, plus budget controls. Use leading-edge models for new development, because they best approximate the AI capability that will be affordable when an application eventually scales.



**Deployed inside your cloud tenancy, SHI’s best-in-class partner solutions from top-tier vendors deliver an intuitive UI, model-agnostic back end, token-level cost controls, and inherited IAM/DLP — so you prototype fast without leaking data or blowing the budget.**

## Prepare to scale

Scalability must be inherent in your AI platform’s design. In the words of SHI’s Shawn Cox, “It has to be scalable because as you drive an application’s adoption, usage will go up – often exponentially.”

### Choose the right hosting model for each workload

Hosting option	Best for	Watch out for
<b>Rent APIs</b> (Azure OpenAI, Anthropic)	Early pilots, bursty traffic.	Unpredictable costs; data privacy concerns.
<b>SaaS copilots</b>	Knowledge worker productivity.	Limited customization; vendor lock-in.
<b>Run your own models</b>	Regulated data; high-volume use.	Significant infrastructure investment.

Design for portability with containerized models and portable vector stores so you can shift workloads as economics or compliance requirements change. Review network architecture to handle streaming responses that may run for tens of seconds.

# Quality and safety: New controls for new capabilities

Generative AI is probabilistic software: the same prompt never guarantees the same answer. This makes automated evaluation of model outputs both more challenging and more important.

SHI AI & Cyber Labs' isolated testbed pairs multiple foundation models with an automated "Reality Check" evaluation harness (PII scrub, bias scan, rubric scoring). Validate prompts and guardrails before anything reaches production.

Validate your AI models with confidence

<sup>6</sup> NIST, "[AI Risk Management Framework](#)"

<sup>7</sup> ISO/IEC 42001:2023: "[Information technology – Artificial intelligence – Management system](#)," 2023

“The core purpose of the SHI AI & Cyber Labs is to help our customers bring their AI use cases to life — use cases that are focused on solving real business problems. We're here to help them iterate quickly and move from concept to production faster than traditional models allow.

- Lee Ziliak, Field Chief Technology Officer at SHI

## Proven playbooks

Don't reinvent the wheel. Leverage proven, publicly available quality and safety control frameworks:

- **NIST AI Risk Management Framework 1.0:** It maps AI risks to practical controls, is voluntary and sector-agnostic (July 2024).<sup>6</sup>
- **ISO/IEC 42001:** The first global management-system standard for AI, using the familiar Plan-Do-Check-Act loop (Dec 2023).<sup>7</sup>

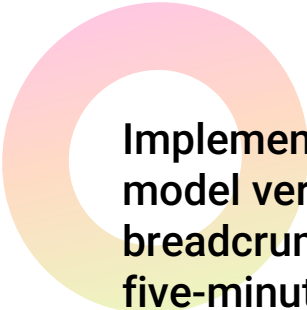
Both frameworks agree that AI assurance is a continuous cycle of *measure » manage » govern » improve*.

## Five frontline safeguards

In addition to whatever industry or business-specific controls your organization requires, make certain you include all of these five:

- 1. Content guardrails:** Block disallowed requests before they reach the model.  
*Example: A “no medical advice” rule stops HIPAA violations at the front door.*
- 2. Privacy shield:** Scrub PII on the way in and out.  
*Example: Regex or ML redaction removes names, e-mails, and account numbers.*
- 3. Reality check:** Compare each answer to a known-good rubric. To automate this evaluation, you must often again use AI, because most generative AI output is inherently variable. This works well in practice! Also, you can – and probably should – use a top-of-the-line model for your evaluation, whether or not you do in production.  
*Example: An evaluation harness scores summaries on coverage, tone, and reading level; bad scores trigger an automatic re-prompt.*

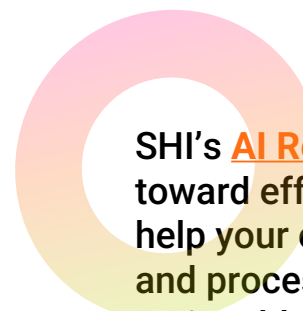
- 4. Human-in-the-loop:** Send high-stakes outputs to a reviewer queue. Depending on the use case, this may be 100 percent or just a sample.  
*Example: Fundraising letters route to advancement staff for one-click approve/edit/reject.*
- 5. Drift watch:** Monitor quality over time to catch degradation early.  
*Example: A nightly job samples 1 percent of answers, re-runs the reality check, and pushes anomalies into the incident channel.*



**Implementation tip: Log everything – the prompt, model version, output, and human decisions. Those breadcrumbs turn audits and post-mortems into a five-minute search instead of a five-week excavation.**

# Implementation approach: From experiment to enterprise scale

Successful AI implementations follow a progressive approach that balances quick wins with sustainable growth. Start small, learn continuously, and scale methodically to build confidence and demonstrate value before larger investments.



SHI's [AI Readiness Workshop](#) is the first step toward effective AI adoption. The assessment will help your organization evaluate current systems and processes, identify potential gaps, and deliver actionable solutions.

Most enterprises spend time in these four stages along their AI journeys. Know what to expect from each.

Phase	Timeframe	Focus
Experiment	Weeks 1–4	Create a secure environment for safe experimentation with low-sensitivity data. Build literacy and identify potential use cases.
Quick win	Weeks 5–12	Develop one high-impact, low-complexity solution for a limited test group. Gather feedback and measure business outcomes.
Scale-out	Months 3–6	Expand successful prototypes while adding 2-3 new use cases. Implement cost monitoring and standardize development patterns.
Process re-engineering	Months 6–12	Integrate AI into core business systems and reimagine workflows around AI capabilities.

Maintain discipline throughout this journey. At quarterly checkpoints, evaluate each use case against defined KPIs. For underperforming initiatives, either fix the underlying issues or retire them and reallocate the resources. This ruthless prioritization keeps investments focused on measurable outcomes.

These fundamentals provide a strong starting point for organizations. Turning these into a fully integrated, enterprise-grade AI capability requires deeper assessment and hands-on execution.

SHI works with organizations to move beyond concepts and point solutions, delivering end-to-end advisory and implementation that connects people, process, technology, and governance. Through structured workshops, technical assessments, and ongoing guidance, SHI helps organizations operationalize AI at scale - aligning strategy to outcomes, managing risk, and accelerating value in a way that is secure, repeatable, and built to evolve.

# Outcomes: From prototypes to a production-ready AI engine

Organizations that invest in strategic AI solutions move beyond scattered experiments and science projects to a production-ready AI engine that employees trust, customers notice, and finance can measure. By unifying data, compute, security, governance, and talent, teams prototype AI use cases in days, scale them in weeks, and refine or augment them as requirements evolve. Clear strategic planning, organizational readiness, and disciplined implementation ensure AI is embedded into core operations, not layered onto legacy systems.

This platform-based approach strengthens readiness across people, process and technology. AI Councils clarify ownership and prioritization, while structured AI expertise enables safe and efficient experimentation within defined guardrails. At the same time, organizations improve data quality, security controls, cost visibility, and operating discipline, creating a stable foundation for scaling AI.


## Business impact: Sustained value at scale

The business impact is both measurable and sustained. High-value use cases are prioritized first. Faster iteration, improved data readiness, and cost-aware model selection then compound productivity gains and prevent runaway spend. Governance and built-in safeguards reduce risk while enabling speed, allowing innovation to advance without compromising trust, regulatory compliance, or brand integrity.

## Why it matters now

AI advantage no longer comes from experimentation alone. Competitive leaders are investing in platforms that turn AI into a repeatable capability.

Building your strategic AI platform



**Bottom line:** A strategic AI platform transforms AI from isolated efforts into a scalable, secure business engine, delivering value today and keeping the organization ready for what comes next.

# About SHI's AI Readiness Workshop

Over the past two years, our experts have successfully deployed generative AI at leading Fortune 1000 organizations. Through this journey, SHI has learned invaluable lessons on achieving success and avoiding major setbacks.

SHI's AI Readiness Workshop is a two-hour, expert-led strategic workshop that benchmarks your technical and business readiness. We evaluate your current systems, data, infrastructure, and business alignment across four critical pillars: security, data, infrastructure, and business.

Our AI experts identify gaps, prioritize quick wins, and deliver a tailored strategic roadmap to enterprise AI adoption. We can assist your teams in AI adoption and increase AI literacy based on our own AI journey and hands-on experience.

**Start Planning and Implementing your AI Journey with us**

Building your strategic AI platform



# About SHI

SHI is not just a global IT solutions provider; we are your AI transformation partner. Our AI Lab infrastructure, in partnership with NVIDIA, is purpose-built for enterprise-grade ideation, experimentation, and adoption. This aligns directly with our strategic approach: Imagine, experiment, adopt, and scale. We empower organizations to imagine what's possible with AI, experiment in secure, scalable environments designed for rapid prototyping and production rollout, and adopt solutions with confidence — backed by our deep expertise in infrastructure, data governance, and enterprise AI.

SHI brings the technical depth and operational rigor needed to succeed. We help you navigate complexity, reduce risk, and unlock real business value.

## Take your AI journey with us.

Connect with an SHI AI expert at [AI@SHI.com](mailto:AI@SHI.com).

Building your strategic AI platform



SHI® is a registered trademark of  
SHI International Corp.