

AI INNOVATION GUIDE | 2026

# Make AI usable in everyday work.

SUHO Co., Ltd. AI practice

AI Adoption / AI Solutions / Local LLM / AI Infrastructure / Physical AI / Operations & Adoption

# An AI partner who knows the infrastructure floor.

Since 2010 SUHO has kept the infrastructure floor running — cloud, data centers, network circuits and round-the-clock monitoring.

On that foundation we support AI end to end: from framing the idea, through GPU and network build-out, to operations and adoption. Our test is not whether AI was installed, but whether it helps the people doing the work.

**15+**

Years in IT ops

**200+**

Clients

**99.9%**

Uptime

**150+**

Projects

**Company**

SUHO Co., Ltd.

**Address**

Kakubari Bldg. 5F, 1-20 Sumiyoshi-cho, Shinjuku-ku, Tokyo 162-0065  
Akebonobashi Sta. (Toei Shinjuku Line) Exit A4, 1 min.

**Founded**

November 12, 2010

**CEO**


Jun ChangSuk

**Business**

Cloud Platform: resale, build, operations, management  
Data Center: rack colocation, circuits, international leased lines  
Product Sales: server, network, storage, security  
Solution: Fortify, NDR, Multi CDN, GameGuard, Microsoft SPLA  
Operation: build, operate, monitor, maintain, incident response  
AI Innovation: AI adoption, Local LLM, AI infrastructure

**Licenses**

Fee-charging employment placement / Worker dispatch / Secondhand dealer  
MIC telecom sales agency notification



## Concept, platform, operations — we stay until the AI keeps running.



### Infrastructure × AI, end to end

01

We have built and run GPU servers, circuits, data centers and cloud ourselves. Beyond picking a model, we prepare the environment that actually runs it — power, installation, network and maintenance included.



### Security-first design

02

We compare on-premises, closed-network and Local LLM options, then design where data sits, who may reach it and how logs are handled — informed by years of shipping security products.



### Start small, stay through adoption

03

From workflow assessment through PoC, production rollout and 24/7 operations. We keep reviewing usage and cost so it grows into something the team keeps using.

# From workflows to infrastructure to the field — six areas that connect AI.

We pick the technology that fits the problem and turn it into something usable — from concept to daily operations.



01 AI PLANNING

## AI Adoption

- Problem framing
- Adoption roadmap
- Validation plan



02 AI SOLUTIONS

## AI Solutions


- Solution selection
- System integration
- Workflow design



03 LOCAL LLM

## Local LLM

- Local LLM adoption
- Internal knowledge
- Access control



04 AI INFRASTRUCTURE

## AI Infrastructure

- GPU & server sizing
- Cloud integration
- Network design



05 PHYSICAL AI

## Physical AI

- Vision
- Edge AI
- Equipment integration



06 OPERATIONS & ADOPTION

## Operations & Adoption

- Monitoring & ops
- User adoption
- Cost & config review



01 AX PLANNING

# AI Adoption

Turn the idea into something the floor can use.

# Not all at once. Proven stage by stage.

Introducing AI is not the goal. We map the workflow, then separate where AI helps from where people should stay in charge. Evaluation criteria are agreed before testing — accuracy, workload and running cost all count.



01 / DISCOVER

## Discover

### What SUHO does

- Map current flow and pain points
- Confirm data location and usage terms
- Compare and prioritize candidate themes

### Deliverables

Issue matrix / candidate themes

02 / DESIGN

## Design

### What SUHO does

- Compare cloud, on-prem and Local LLM
- Design permissions and review flow
- Set criteria, team and schedule

### Deliverables

Architecture / PoC plan / criteria

03 / VALIDATE

## Validate

### What SUHO does

- Prepare environment, run scenarios
- Assess accuracy, latency, review load
- Analyze gaps, set go-live conditions

### Deliverables

Results / fixes / go-live evidence

04 / ADOPT

## Adopt

### What SUHO does

- Roll out to production in stages
- Document procedures and rules
- Review usage and cost continuously

### Deliverables

Rollout plan / runbook / ownership

\* Scope, deliverables and duration are defined case by case, based on the target workflow and data.

# Find the opening in everyday work.

Topics clients raise most often. We scope them to your environment, keep automation bounded, and route anything hard to judge back to a person.



## Cut document checking and re-entry

Pull the fields you need from forms and ledgers, let a person confirm, then write to your existing system. We look at exception handling and approval flow, not just read accuracy.

Document

Extract fields

Human check

To your system



## Make internal information findable


Help people reach policies and manuals. We cover how content is updated, who may view it, and how the source is shown — so the answer can be trusted at work.

Question

Search

Cite source

Answer



## Rework repetitive triage

Support the routine calls — classifying inquiries, suggesting who should handle them — and hand the hard ones to a person.

Intake

Suggested class

To owner

Handled



02 AI SOLUTIONS / LOCAL LLM

# AI Solutions & Local LLM

Turn internal knowledge into internal capability.

# Generative AI — your data stays in.

We deploy LLMs on-premises or in closed networks to match your security requirements — designing everything from GPU server sizing and installation through to operations.



## A fit when

You handle customer data

R&D data must stay in

External AI is off-limits

API cost must be fixed

### Using a cloud LLM as-is

- Confidential data leaves for an outside service
- Metered billing grows with every new user
- Internal rules or audits may forbid it outright

### With a Local LLM

- **Data stays inside your network**
- **Flat-rate GPU capacity makes budgeting predictable**
- **Access rights and logs stay under your control**

## What we deliver



### Runs fully on-premises

Inference with no external traffic — GPU selection, facilities and maintenance included



### Open-source model selection

Llama, Qwen, Gemma and others compared on fit, accuracy and speed



### Internal knowledge search (RAG)

A search layer that safely reaches policies, manuals and past documents




### Permissions & operations


Access control by team and role, usage logs and update procedures

## Connected to your systems, built into the work.

For document work, information retrieval and daily repetition. We select the right AI solution and design its integration and workflow with your existing systems.

What AI handles and where a person checks are drawn as clear lines, so the team can trust what it uses.





03 AI INFRASTRUCTURE / PHYSICAL AI

# AI Infrastructure & Physical AI

Run the AI — and everything after it.

# Design the whole stack, not just the GPU.

Which model, how many users, how the data is handled. We build the configuration from those constraints, with operations and growth already in view.



01 / CLOUD

## Cloud

Cloud compute that follows your pilot and your changing load.

**Good fit for** Pilots / variable workloads / shared multi-site use

**Points to check** GPU regions and quotas, runtime hours, egress fees, retention terms



02 / ON-PREMISES

## On-premises

A dedicated environment in your office or DC, run under your control.

**Good fit for** Local LLM / internal data / sustained dedicated use

**Points to check** Power, cooling, space, capex, maintenance, refresh, egress control




03 / HYBRID

## Hybrid

Split what stays inside from what runs in cloud, joined into one workflow.

**Good fit for** Reuse existing kit / phased migration / burst to cloud

**Points to check** Transferable data, latency, line cost, federated auth, fault isolation



## Four layers behind an AI platform



### GPU & servers

Sizing, configuration, runtime



### Network & access

On-prem to cloud links and auth



### Data & storage

Location, retention, backup



### Monitoring & ops

Monitoring, runbooks, cost review

## 05 / PHYSICAL AI


## Bring AI to the floor.

Use camera and sensor data to improve work on site. Vision, edge AI and equipment integration — configured and proven for your environment.



### Cameras & sensors

Collect site video and readings



### Edge AI

Recognize and judge on site



### Equipment & workflow

Alerts, records, machine control

## 06 / OPERATIONS &amp; ADOPTION

## We stay past go-live.

Go-live is where operations begin. We extend the same round-the-clock monitoring we run for infrastructure, watching uptime, usage and cost while turning feedback from the floor into improvements.



### Monitoring & ops

Uptime and usage monitoring, incident escalation and recovery, update and maintenance runbooks



### User adoption

Written procedures, user rollout and helpdesk, gathering feedback from the floor



### Cost & config review

Track usage and cost, revisit capacity, configuration and model choice

# A first conversation, or the full build.

Too often the words AI and DX arrive before any plan does. We match the engagement to your stage — from a single advisory session to running the platform for you.



## SPOT

### Advisory

Start by framing where you are

- Frame the problems, list candidate use cases
- Advisory on technology and architecture
- Draft an adoption roadmap



## PROJECT

### PoC & rollout

Prove it, then take it live

- Environment setup through scenario runs and scoring
- Cloud, on-prem or Local LLM architecture and build
- Integration with existing systems and workflow build



## PARTNER

### Managed partnership

Hand us the running of it

- Phased rollout into production
- 24/7 monitoring, operations and maintenance
- Continuous review of usage and cost

\* Cost and duration depend on how many workflows, the state of your data, integrations and evaluation scope. We scope and quote after an initial session.

# We use it ourselves first — AI already changed how we work.

Before proposing anything, we build and run AI and cloud on our own operations. What works and what trips you up goes straight into your project.

## In-house

### Automated quotations

LLM × accounting SaaS

Supplier quote PDFs are read by AI, line items and prices extracted, and a quotation created in the accounting cloud. A human confirms before saving, cutting transcription errors.




## In-house

### Sales & billing AI assistant

Natural language → SQL → answer

A portal unifying pipeline and billing data, with an assistant that turns a question into SQL and answers it. Read-only, so it stays safe.



## In-house

### Email work console

Summarize, classify, turn into tasks

Work email is summarized and classified by AI, with replies and follow-ups tracked as tasks. Runs serverless in the cloud so nothing is missed.



## Client work

### Site Wi-Fi analytics

Measured data, visualized

Built for the Japanese arm of an automotive manufacturer: collects live wireless AP telemetry and visualizes coverage. No install needed on site.



\* Screens shown use demo data.

LET'S START A CONVERSATION

# Start with one problem.

'This step eats our time.' 'I want fewer checks here.'  
No formal plan needed — just tell us where it hurts.



[contact@suhojapan.com](mailto:contact@suhojapan.com)



[suhojapan.com](https://suhojapan.com)



1-20 Sumiyoshi-cho, Shinjuku-ku, Tokyo 162-0065

SUHO Co., Ltd.