

Realigns Inc- Research and Innovations Artificial Intelligence

RINC AI Local V2

Private AI from everyday PCs to industrial GPU systems



RINC AI Local V2

User Deployment Brochure



Local, Offline-capable, and Private




Model-flexible and Enterprise-ready



Your Data. Your AI. Your Infrastructure.



Scales from compact local models to large industrial deployments



RINC at a Glance

One local AI workstation. Multiple workflows.

Core Operating Principle: Use the correct module and model for the workload. Importing a model does not prove that the model is appropriate for the machine or task.

01

Private by Design

Workloads remain on the user machine or approved on-premises host.

02

Model-Flexible

Use compatible models sized for the machine, task, license, and quality.

03

Deployment-Flexible

Start on localhost, then scale to private LAN or multi-host patterns.

User Workflows

What can one RINC workstation do?

Important Robotics Boundary: RINC interprets. Customer software validates and executes. Physical safety logic remains customer-controlled.

01

Chat Studio & Documents

Drafting, summarization, and business support via local text models, plus private file-grounded retrieval and document Q&A.

02

Research & Coding

Hybrid research when external information is required, alongside local code drafting, bug analysis, and refactoring.

03

Image & Voice

Dedicated local image generation with separate runtimes, and local speech recognition and synthesis.

04

PPT, API & Robotics

Structured presentation drafting, authenticated internal API access, and governed interpretation for robotics and vision.

Single-User Deployment

Start on the computer you already own.

System memory	Typical model class	Acceleration	Best fit
4 GB*	~300M-1B Q4	CPU	Demo / basic text
8 GB	0.5B-3B Q4	CPU / iGPU / Metal	Compact chat / light vision
16 GB	3B-7B Q4	CPU + GPU/Metal	Daily local AI
32 GB	7B-14B	GPU/Metal	Higher quality / larger context
64 GB+	14B-32B+	High-memory GPU	Advanced local workflows


01 Review **RAM**, CPU/GPU, and storage

02 Import/select approved **model**

03 Start local **runtime** and verify output before impact use

Model Scale Ladder

From ~300M to 120B+ models.




Ultra-Light (4-8 GB)

~300M-1B: Basic chat, routing, simple extraction.



Compact (8 GB+)

0.5B-3B: Everyday local assistant, light vision.



Balanced (16 GB+)

3B-7B: Documents, coding, research synthesis.



Advanced (32 GB+)

7B-14B: Harder reasoning, larger context.



Large Local (64 GB+)

14B-32B+: Enterprise workstation, workflows.



Industrial & Up

70B-120B+: High quality, concurrent use.

Industrial and Enterprise Deployment

Scale the host, not the privacy boundary.

Deployment	What it does	Infrastructure focus
Single-room host	One AI workstation serves office LAN	Stable private IP, firewall, authenticated clients
Multi-floor routed host	One host serves floors/subnets	VLANs, inter-VLAN routing, approved CIDRs
Multi-host department	Separate hosts handle workloads	Workload isolation, independent models, IT governance

Private Connectivity

Internal services without public internet exposure.

Security Boundary: Do not expose RINC API to the public internet. Security Console does not replace firewall, EDR, or identity controls.

01 **Local API**
Internal SaaS and department workflows over localhost or private LAN.

02 **Network Controls**
Private subnets, VLANs, firewall restrictions, and least privilege.

03 **Robotics API**
Governed text, vision, and read-only telemetry for industrial software.

04 **Security Boundary**
Do not expose RINC API to the public internet or replace core controls.



Choose Your Deployment

Three practical user scenarios.

01

Personal / Professional

Laptop/desktop, 4–16 GB RAM, ~300M–7B models, CPU acceleration.

02

Team / Department

Workstation or AI server, 32–64 GB RAM, 7B–32B+ models.

03

Industrial / Enterprise

Multi-GPU systems, 96 GB+ RAM, 70B–120B+ models, private VLAN.

Before You Deploy

A simple user checklist.

A close-up photograph of various electronic components, including resistors and capacitors, scattered on a dark, textured workbench. The lighting is warm and focused on the components.

Hardware

Confirm RAM, acceleration, storage, and thermal capacity.

A person with short dark hair is shown in profile, looking at a large computer monitor. The monitor displays a complex data dashboard with multiple charts, graphs, and a map. The person's hand is pointing at the screen.

Models

Record approved source, license, and quantization.

A wide-angle photograph of a modern warehouse interior. High industrial shelving units are filled with boxes. A small, white, box-like robot is visible in the aisle, and green lights are mounted on the shelves.

Runtime

Verify Ready status before dependent workflows.

A photograph of a small, green, leafy plant in a white ceramic pot. The plant is sitting on a white desk next to some papers and a pen. The lighting is soft and natural.

Privacy & Network

Distinguish local services and use secure network routes.

A photograph of a large cargo ship at night, illuminated by its own lights. The ship is filled with colorful shipping containers in shades of orange, blue, and red. The ship is docked at a pier.


Security & Capacity

Protect API keys and load-test host combinations.

A photograph of a person wearing a full-body protective suit, including a hood and gloves. The person is holding a tool, possibly a screwdriver, and is working on something. The suit has the 3M logo on it.

Human Review

Maintain responsibility for professional judgment.



RINC AI Local V2

From a normal computer to industrial AI infrastructure.



Start with the hardware you have



Choose the model that fits



Scale when workload demands it



Your Data. Your AI. Your Infra.