


Your test content stays in your network.

The only test content that can leave is the inference payload: the current screenshot and step text. With a self-hosted model, not even that.



Where inference runs is your choice

Default

AskUI hub

Inference routed and billed through your workspace. Customer data and screenshots never train models.

What most IT reviews approve

Your cloud tenancy

Bring your own model: your Anthropic account, Azure AI Foundry, or a gateway behind Microsoft Entra.

Air-gapped

On-premise inference

A self-hosted vision model with vLLM or SGLang. Offline licence key, no hub, no AskUI account.

The blast radius is the toolbox

Screen and input	Come with the device; the agent can do what the logged-in user can do.
Beyond the UI	Shell, SQL and file tools are off until enabled, and enabling is recorded in the project (utils/tools.json).
Custom tools	C# code in your repository, reviewed like any other code.
Elevation	Administrative apps only if explicitly allowed on the AgentOS service.
Credentials	Named secrets, encrypted locally, never shown to the model, redacted from reports and logs.

Network, telemetry, compliance

Inbound	Only the AgentOS gRPC port on target machines (23000 app-spawned, 26000 service). Restrict it to the machines that drive tests.
Outbound	HTTPS to the model provider you chose, and to the hub for sign-in and usage (not needed in licence-key mode).
Telemetry	OpenTelemetry off by default; if enabled, to your own collector.
Certification	ISO/IEC 27001:2022. GDPR compliant. DPA available.
Company	askui GmbH, Karlsruhe, Germany. HRB 740125.