Attestation and confidential computing

Service terms and information. The customer’s executed agreement, order, and schedules govern.

1. Service technical specification

LoricAI protects documents, keys, and processing in its confidential computing perimeter. This specification, version 2026-11-12, describes platforms, data flow, and verification. Signed evidence for the assigned configuration allows the IT contact to check the environment; this description does not replace that evidence.

2. Hardware platforms

LoricAI infrastructure offers hardware enclaves on NVIDIA RTX PRO 6000 Blackwell Server Edition and NVIDIA B200. AMD CPUs and NVIDIA GPUs contribute to verification of the protected environment. The assigned profile identifies the platform, protection mode, and customer configuration. The service configuration identifies CPU, GPUs, firmware, drivers, and confidential-computing mode. A GPU model name alone does not prove protected mode is enabled: CPU evidence and evidence for every assigned GPU are checked.

3. Model and location

The AI model is Qwen 3.8 Flash Next.
Qwen is an open model developed by Alibaba.
We run it only on our protected servers.
Alibaba receives neither your documents nor your questions. Documents are processed in Finland, in the European Union, on Verda servers.
Keys are managed separately through Google Cloud KMS. Cloudflare - network; encrypted R2 storage in the European Union.
Backblaze - encrypted B2 backups in the European Union. The customer configuration sheet associates AI model/version, authorized software, and service identity with the execution profile.

4. Encrypted document, protected processing

Device encryption protects transfer and storage. To process a request, data is decrypted only inside the authorized enclave where the AI model runs. The answer returns to your workspace over a protected channel. Protected memory and channels form part of the verified configuration. Content is not sent to external AI APIs or used for training. Each customer has a dedicated key, managed separately from documents through the key-management service. The key-management service does not receive documents. The release policy authorizes the protected environment after integrity checks; credentials, roles, and access logs govern key operations.

5. Step 1: identify the profile

Obtain service identity, CPU, assigned GPUs, model version, approved software, reference values, and key-release policy. Check the profile belongs to the customer workspace.

6. Step 2: gather evidence

Obtain CPU and GPU attestations, certificate chains, and fresh proof tied to the session. Old or unrelated evidence does not describe the current session. Before sending a document, your app asks the chip for proof signed by AMD and NVIDIA.
The proof shows that the protected area is sealed and runs only the program we published.
If the proof does not check out, the document is not sent.

7. Step 3: verify signature and configuration

Use official AMD and NVIDIA verifiers compatible with the profile. Check signatures, validity, certificate revocation, measurements, firmware, protected mode, and software identity against authorized values. Repeat for every GPU. GPU attestation does not replace CPU attestation or vice versa.

8. Step 4: check key-release policy

Confirm invalid, expired, or policy-mismatched evidence prevents key release. Use the agreed verification path with demonstration data, without production documents. Record result, time, configuration, and verifier versions in the customer register.

9. Protection and continuity

Attestation checks environment and policy, not answer quality, application permissions, or availability. We keep encrypted backup copies for 5 years.
They are protected by your key: without it, nobody can read them. Authority requests and transfers are governed separately from processing location and hardware evidence.

10. Operational references

LoricAI guide: https://loricai.io/en/technical-verification/
AMD SEV: https://www.amd.com/en/developer/sev.html
NVIDIA Attestation: https://docs.nvidia.com/attestation/
NVIDIA platforms: https://docs.nvidia.com/confidential-computing/containers/v1.1.0/supported-platforms
Technical contact: support@loricai.io.