

NAGE

COMPANY INTRODUCTION

Knowledge you can inspect



Meet Nage

AI infrastructure for source-aware outputs and verifiable records.

KLM / SEDIM / Nage AI OS / Craton

For enterprise teams and technology partners

A clearer basis for AI decisions

When an AI output enters a business process, the next question is often about its basis.

The operational need

A policy answer, an internal recommendation or a research summary may combine retrieved documents, user input and a model's prior knowledge. Teams need to examine the sources presented, the versions in use and the signals attached to the result.

The Nage approach

Nage develops KLM, a candidate standard for recording the declared basis of AI outputs, and operates its reference implementation. SEDIM supplies the architecture; Nage AI OS carries the records across services. The aim is a trust anchor: a shared record that another party can inspect within a stated scope.

The value to your team

Review an answer with its source identity, version and declared signal status. Keep the record available as work moves between an application, a reviewer and an assurance team. Start with one workflow and agree what evidence is needed for adoption.

Record verification establishes specified properties of a record. It does not guarantee that an answer is true.

One ecosystem, distinct roles

A candidate standard, its reference implementation and the services that carry the record.

KLM

The candidate standard

Knowledge Layers Model defines a shared record vocabulary. The public charter describes the candidate standard; the supplied internal specification develops its conformance requirements.

SEDIM

The architecture

Separates the base-model role from managed knowledge layers and connects outputs to source and composition records.

Nage AI OS

The operating environment

Carries source traces and records across knowledge services, delivery interfaces and agent interactions.

Craton / Craton Pro

The base-model programme


A from-scratch base and its larger class, designed to reduce knowledge that cannot be attributed to managed layers. Both remain in development.

The current SEDIM service uses an existing open-weight base. Trained VARVE weight injection is a development path, rather than the current production inference path.

The Nage architecture

KLM defines the record contract; Nage operates the reference implementation.

KLM / CANDIDATE STANDARD
Shared record vocabulary, signal status and conformance requirements



Nage AI OS carries the record across services
Application / reviewer / assurance tooling

Current base Qwen3-8B in the FACIES slot. Current grounding and record service.	Development programme Craton / Craton Pro and trained VARVE weight composition.
---	---

Conceptual service flow, not a causal execution proof. Current inference uses grounding with Qwen3-8B; trained VARVE weight injection and Craton remain development paths.

What we can explore together

A focused conversation can connect your workflow to a testable evaluation.

A useful starting point

Bring one knowledge-intensive workflow, an example of the sources it depends on and the review decision you need to support. We can discuss the record, integration boundary and evidence an evaluation should produce.

ENTERPRISE TEAMS

Source-aware knowledge workflows

Review an answer alongside the evidence offered for it. Explore how source versions, missing signals and retirement events could fit your internal review process.

TECHNOLOGY PARTNERS

Records that travel across systems

Discuss how applications, infrastructure providers and assurance tools could exchange KLM records while preserving source identity and signal status.

Use cases describe evaluation opportunities. They are not customer references or claims of regulatory approval.

Let's discuss your use case

Meet the Nage team during the conference to explore source-aware AI and independently checkable records.

In a product conversation

Explore the current SEDIM service, the information an output record carries and which checks a reviewer can perform.

In a technical conversation

Examine record integrity, disclosure requirements, knowledge lifecycle operations and the distinction between a source reference and evidential support.

In a partnership conversation

Discuss application integration, domain knowledge and verification tooling. Craton and Craton Pro open a further conversation about the base-model layer.

Reply to the accompanying invitation email to arrange a conversation.

nage.ai

For event dates, location and meeting arrangements, please refer to the accompanying invitation email.