GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

Georg-August-Universität Göttingen / GWDG

HPS NHR-NORD@GÖTTINGEN

https://hps.vi4io.org/

Knowledge-Guided HPC Decision Intelligence: Executable Governance for Trustworthy AI-Driven Decision Workflows

Aasish Kumar Sharma and Julian M. Kunkel

Presenter: Aasish Kumar Sharma

15.09.2026

Table of contents

1 Introduction & Background

2 Problem

3 Proposed Solution

4 Takeaway

Introduction: This Contribute Talk Is About

What: An Issue?

Today, AI assist everywhere, including decision making.

Introduction: This Contribute Talk Is About

What: An Issue?

Today, AI assist everywhere, including decision making.

Why This Is An Issue?

Decisions need evidence, explainability, & accountability.

Introduction: This Contribute Talk Is About

What: An Issue?

Today, AI assist everywhere, including decision making.

Why This Is An Issue?

Decisions need evidence, explainability, & accountability.

How?: Design

System, validate each decision's evidence.

Introduction: This Contribute Talk Is About

What: An Issue?

Today, AI assist everywhere, including decision making.

Why This Is An Issue?

Decisions need evidence, explainability, & accountability.

How?: Design

System, validate each decision's evidence.

That Means

AI-assisted HPC decisions need executable governance system by design, that can validate decision's evidence on the spot.

Introduction: This Contribute Talk Is About

What: An Issue?

Today, AI assist everywhere, including decision making.

Why This Is An Issue?

Decisions need evidence, explainability, & accountability.

How?: Design

System, validate each decision's evidence.

That Means

AI-assisted HPC decisions need executable governance system by design, that can validate decision's evidence on the spot.

Goal: How to make real time executable governance possible?

AI In HPC Decision Making: A Scenario

Policy: HPC Decision Setting

- GPU hours are limited and expensive.
- Queues must be fair and transparent.
- Single allocation must not exceed 20%.
- AI can assist in scheduling decisions.

AI In HPC Decision Making: A Scenario

Policy: HPC Decision Setting

- GPU hours are limited and expensive.
- Queues must be fair and transparent.
- Single allocation must not exceed 20%.
- AI can assist in scheduling decisions.

Governance concerns

- Was it logged?
- Was it explainable?
- Was it traceable?
- Was it fair under the active policy?

AI In HPC Decision Making: A Scenario

Policy: HPC Decision Setting

- GPU hours are limited and expensive.
- Queues must be fair and transparent.
- Single allocation must not exceed 20%.
- AI can assist in scheduling decisions.

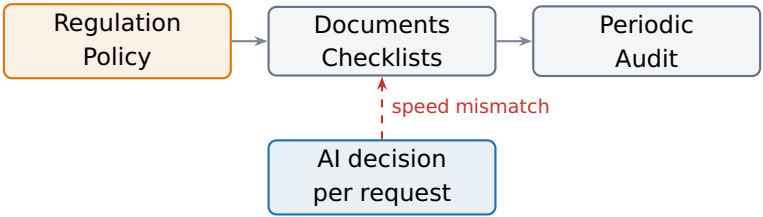
Governance concerns

- Was it logged?
- Was it explainable?
- Was it traceable?
- Was it fair under the active policy?

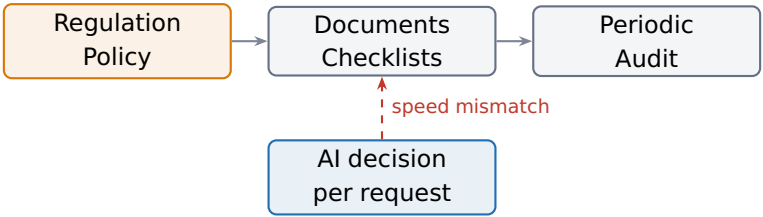
The main question!?

How to generally make the governance executable & evidence verifiable?

The Problem: Governance Is Still Mostly Documentational



The Problem: Governance Is Still Mostly Documentational



Core gap

- **A document checklist (obligations) claim compliance.**
- **But it cannot verify every decision's evidence while executing a decision.**

OKB paper: compliance remains documentation-centric and evidence-light.

State of the Art & Missing Layers

Approach	Provides	Missing layer
Regulation	Obligations	Executable checks
Documentation	Audit artefacts	Per-decision validation
XAI methods	Model explanations	Governance verdicts
Policy-as-code	Rule execution	Semantic evidence graphs
Compliance ontologies	Structured meaning	Runtime profile validation

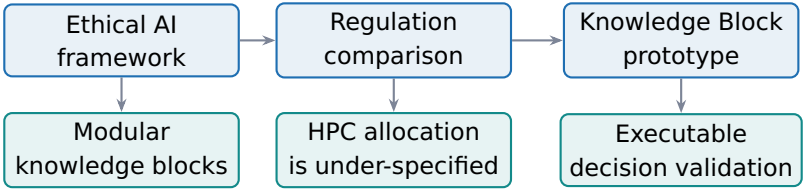
State of the Art & Missing Layers

Approach	Provides	Missing layer
Regulation	Obligations	Executable checks
Documentation	Audit artefacts	Per-decision validation
XAI methods	Model explanations	Governance verdicts
Policy-as-code	Rule execution	Semantic evidence graphs
Compliance ontologies	Structured meaning	Runtime profile validation

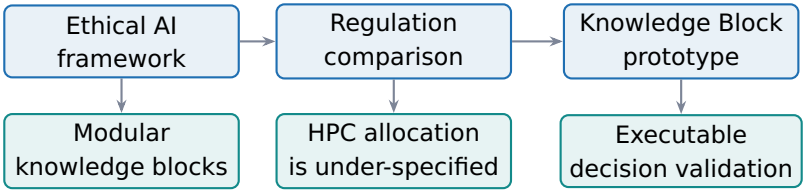
Objective

To target executable, evidence-grounded governance for AI-assisted HPC decisions.

How The Work Can Evolve



How The Work Can Evolve



Research question

Can reviewed obligations become deterministic checks over decision evidence?

Ethical AI COMPSAC 2025, Global AI Regulations COMPSAC 2026, OKB COMPSAC 2026.

Possible Solution: Ontological Knowledge Blocks (OKB)

Definition

An OKB, a machine readable set of knowledge blocks that turns one reviewed obligation into one executable validation unit.

Possible Solution: Ontological Knowledge Blocks (OKB)

Definition

An OKB, a machine readable set of knowledge blocks that turns one reviewed obligation into one executable validation unit.

Formal abstraction

$$KB = \langle O, C, V, E, P \rangle$$

O: obligation C: concepts V: validation E: evidence P: provenance

Possible Solution: Ontological Knowledge Blocks (OKB)

Definition

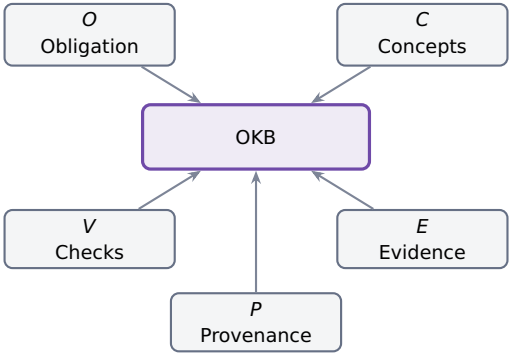
An OKB, a machine readable set of knowledge blocks that turns one reviewed obligation into one executable validation unit.

Formal abstraction

$$KB = \langle O, C, V, E, P \rangle$$

O: obligation C: concepts V: validation E: evidence P: provenance

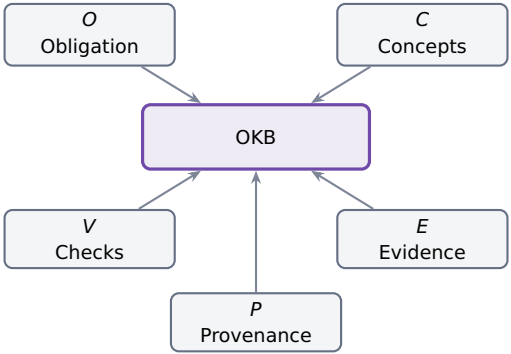
What Exactly Is an OKB?



Concise Meaning

- It is a governance rule package.
- It states what must be checked.
- It states which evidence is needed.
- It executes as SHACL validation.
- It make a machine-readable verdict.

What Exactly Is an OKB?



Concise Meaning

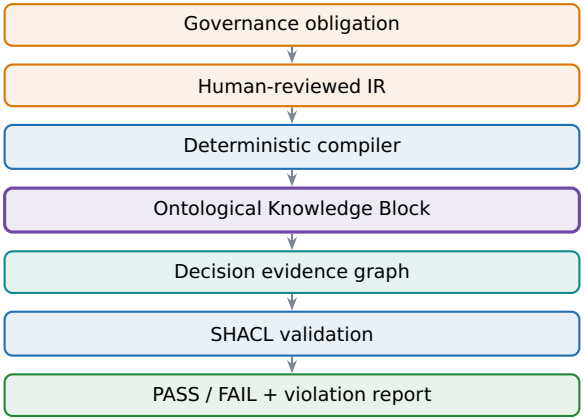
- It is a governance rule package.
- It states what must be checked.
- It states which evidence is needed.
- It executes as SHACL validation.
- It make a machine-readable verdict.

Example

Fairness OKB checks whether GPU-hour disparity exceeds the approved threshold.

Governance • **Mechanism** • **OKB** • **Evidence** • **Pass** • **Violation**

How It Works: From Obligation to Verdict



Human interpretation remains human. Execution becomes deterministic.

OKB Provides: Governance Profiles Change Obligations

Accountability

- Usage log.
- Timestamp.
- Model trace.

OKB Provides: Governance Profiles Change Obligations

Accountability

- Usage log.
- Timestamp.
- Model trace.

Fairness

- Explanation.
- Allocation values.
- Disparity threshold.

OKB Provides: Governance Profiles Change Obligations

Accountability

- Usage log.
- Timestamp.
- Model trace.

Fairness

- Explanation.
- Allocation values.
- Disparity threshold.

Combined: Transparent

- Accountability.
- Fairness.
- Same validator.

OKB Provides: Governance Profiles Change Obligations

Accountability

- Usage log.
- Timestamp.
- Model trace.

Fairness

- Explanation.
- Allocation values.
- Disparity threshold.

Combined: Transparent

- Accountability.
- Fairness.
- Same validator.

Important Property

- **Governance changes by profile selection.**
- **Does not require rewriting the AI service.**

Finding 1: Profiles Detect Intended Failures

Evidence issue	Accountability	Fairness	Combined
Missing explanation	✔	✘	✘
Missing model artefact	✘	✔	✘
Excess disparity	✔	✘	✘

Conformant decision

120 vs. 110 GPU-h

$r = 0.083 < 0.20$

✔ **PASS**

Violating decision

200 vs. 50 GPU-h

$r = 0.75 > 0.20$

✘ **FAIL**

Finding 1: Profiles Detect Intended Failures

Evidence issue	Accountability	Fairness	Combined
Missing explanation	✔	✘	✘
Missing model artefact	✘	✔	✘
Excess disparity	✔	✘	✘

Conformant decision

120 vs. 110 GPU-h

$r = 0.083 < 0.20$

✔ **PASS**

Violating decision

200 vs. 50 GPU-h

$r = 0.75 > 0.20$

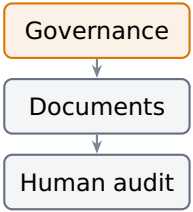
✘ **FAIL**

Overall: Evidence-Profiles-Satisfiability

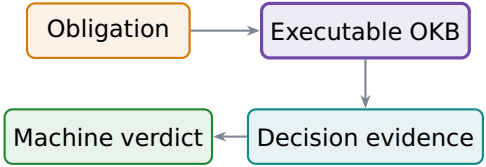
Each profile fires on the failure class it is designed to govern.

What We Contribute

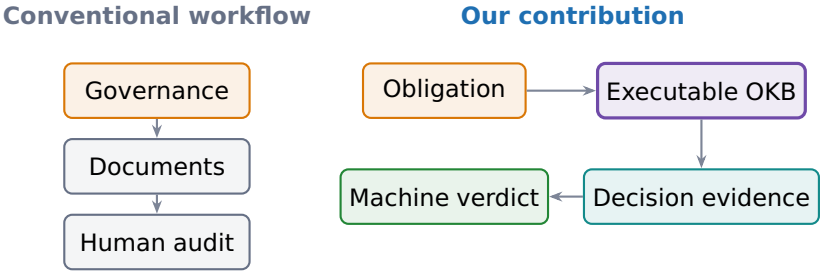
Conventional workflow



Our contribution



What We Contribute



Main contribution

Governance obligations become machine executable checks over decision evidence.

Contributed Papers: Supporting Publications

1. Ethical AI: A. K. Sharma, D. Kyosev, and J. Kunkel, “Ethical AI: Towards Defining a Collective Evaluation Framework”, IEEE COMPSAC 2025, DOI: 10.1109/COMPSAC65507.2025.00344.

2. Global AI Regulations: A. K. Sharma et al., “Global AI Regulations for FAIR and Ethics in High-Risk Use Cases: A Comparative Review”, IEEE COMPSAC 2026, DOI: 10.1109/COMPSAC69091.2026.00092.

3. OKB Paper: A. K. Sharma and J. M. Kunkel, “Ontological Knowledge Blocks: Executable Compliance and Profile-Based Validation for Trustworthy AI Systems”, IEEE COMPSAC 2026, DOI: 10.1109/COMPSAC69091.2026.00379.

Takeaway

Final message

A good HPC decision should not only be optimal;

it should also be evidenced, explainable, and verifiably governed.

Aasish Kumar Sharma · Julian M. Kunkel

Georg-August-Universität Göttingen / GWDG

Questions & Discussion

Contact: aasish-kumar.sharma@gwdg.de