



Human Authority for Agentic AI

Capability is not authority.

A narrow architectural thesis: an enterprise agent may be technically capable of an action and permitted by runtime policy, yet still lack a valid human mandate for the consequential business effect.



PROPOSED INTEGRATION BOUNDARY



Integration principles

- Technical permission must not silently create human authority.
- Unknown or ambiguous authority fails closed; the agent cannot widen its own grant.
- A one-shot grant is consumed after use; replay requires new authority.
- The authority decision produces evidence that can be reconstructed later.

The smallest credible NVIDIA x LHINOS proof

One NVIDIA-based agent. One consequential action. One independent LHINOS authority decision. No replacement of NVIDIA runtime security or guardrails.

VALIDATION CASES

01 No human grant

DENY

Capability + technical PASS, but authority absent.

02 Explicit handback

HUMAN_CONFIRMATION_REQUIRED

Return control to a human; do not infer authority.

03 Bounded one-shot grant

ALLOW

Actor/action/target/purpose/limit/time/evidence match.

04 Reuse after consumption

DENY

Action budget = 0; silent replay is prevented.

05 NVIDIA policy blocks

BLOCKED

Human authority never overrides technical policy.

06 Expired / mismatched grant

DENY

Scope/target/purpose/threshold/validity mismatch fails closed.

MINIMAL CONTRACT

AuthorityRequest

```
{
  "actor": "agent:enterprise-ops",
  "action": "customer_record.update",
  "target": "crm:case-1842",
  "purpose": "approved_case_resolution",
  "requested_budget": 1
}
```

AuthorityDecision

```
{
  "decision": "ALLOW",
  "grant": { "action_budget": 1,
  "valid_until": "...",
  "self_widening": false },
  "evidence_receipt": "..."
}
```

A VALID FALSIFIER

What would reduce or eliminate the LHINOS fit?

If NVIDIA already enforces equivalent end-to-end human-authority semantics - named human source, actor/action/purpose/target constraints, bounded validity and budget, non-self-widening and reconstructable evidence - the LHINOS thesis must narrow or stop.



LIVE DEALROOM

Navigator + Authority
Proof + ChatTHINOS
+ Translator

Requested NVIDIA outcome - asynchronous technical assessment

<https://lhinos-hi.com/nvidia/>

Fit worth testing | Likely overlap | Not relevant

info@lhinos-hi.com

NVIDIA PUBLIC REFERENCES

NeMo Guardrails Tool Calling: docs.nvidia.com/nemo/guardrails/configure-guardrails/guardrail-catalog/tool-calling
OpenShell Security Best Practices: docs.nvidia.com/openshell/dev/security/best-practices