OpsChainAI / APEX

# Governed AI execution, tested in Azure.

## Why this matters

An agent can have a legitimate identity and still attempt an action outside its task. Enterprise controls must determine which resource it may access, bind authorization to the actual request, and preserve limits when infrastructure restarts.

## What APEX demonstrated

On September 24, 2026, the isolated APEX Azure test deployment passed five SQL checks: approved access returned the expected record; a restricted customer was denied; altered arguments were rejected; a completed execution could not be reused; and the original signed request remained valid after the tampering attempt.

## What was connected

The gateway used Microsoft Entra authentication, signed execution authorization, an exact resource policy, and a dedicated managed identity to access Azure SQL. The database held synthetic records. Its identity could read both test customers, while the caller was permitted to read only one.

## Continuity of enforcement

Separate live tests replaced the container and confirmed that replay and action-budget state persisted in Azure Table Storage. Two concurrency tests each admitted one of six competing requests. Audit checks separately rejected unlocked retention and accepted a blob with locked version-level retention.

Evidence baseline: gateway commit f990a45; SQL test revision 0000006. Results are controlled validation, not production certification.

# A practical pilot with a defined boundary.

## Proposed engagement

Select one enterprise workflow and a synthetic dataset. Agree on the permitted identity, operation, resource, and action budget. Demonstrate normal execution and deliberate deviations, then provide a results report and deployment-specific findings.

## What the evidence does not establish

The tests do not reproduce the OpenAI/Hugging Face intrusion or establish complete agent containment. Artifactory integration, direct network bypass, host compromise, and trusted signing-key theft were not validated. Independent administrative custody of the audit store has not been demonstrated.

## Next engineering milestone

Isolate the agent workload from backend credentials and signing material, and test whether it can reach protected services outside the gateway. The enforcement claim depends on requests passing through the governed path.

## Discuss a scoped validation

Contact paul@npmit.com or visit opschainai.com. Pilot scope, deployment requirements, timing, and acceptance criteria should be agreed before work begins.

## Incident context

OpenAI: The Hugging Face incident and the road ahead. Hugging Face: Anatomy of a Frontier Lab Agent Intrusion. These reports motivate the scenarios; APEX test results are a separate body of evidence.

OpenAI incident report

Hugging Face technical timeline