

Sovereign Control Plane for Enterprise AI Agents

Governing action and preserving institutional memory

Russ James
Chief Architect, COREONYX

QURii patent awarded. Additional sovereign agent technology Patent Pending.

SEPTEMBER 2026 | ADDITIONAL TECHNOLOGY PATENT PENDING

Own the authority.
Choose the intelligence.
Keep the institutional
memory.



Executive perspective

The enterprise should retain authority over what its agents do and the knowledge their work creates.

A team needs an answer, a system update, or a completed handoff. An AI agent can help move that work forward. But the person accountable for the outcome must still know which information was used, whose permissions were exercised, what was approved, and what actually happened.

That is the operating problem QURii addresses. It places a durable, policy-aware control plane between model reasoning and operational systems. Models propose structured actions. The platform governs their path to execution through declared capabilities, explicit identities, policy decisions, and human approval where required. A correlated record preserves the available context and results.

This paper explains the sovereign agent platform architecture behind QURii, its seven pillars, and the practical questions an enterprise should use to evaluate it. The audience includes CIOs, CISOs, enterprise architects, platform leaders, mission owners, and technical acquisition teams in government and commercial organizations.

Three decisions for leadership

Decision	What the enterprise should establish
Authority	Who may authorize each action, which systems it can affect, and how execution is constrained.
Continuity	How work resumes after interruption and how the next authorized person understands its history.
Choice	How models, infrastructure, and integrations can evolve while the enterprise governs its operating knowledge.

QURii’s proposition is the combination of durable execution, mediated action, explicit identity, deterministic policy and isolation, provider choice, history ownership, and customer authorship. No single pillar is claimed as unique in isolation. The value lies in how they work together around a consequential workflow.

Technical posture: September 2026. Current implementation and planned extensions are separated on pages 10–11. Architecture and deployment-dependent capabilities are not guarantees of security, compliance, or measured business results.

Reading paths: executive argument 2–3 • architecture and controls 4–7 • workflow and evaluation 8–9 • implementation posture and references 10–12.

What sovereignty means in practice

In this paper, sovereignty means meaningful enterprise control over the conditions under which agents operate. It includes the data they can use, the infrastructure they run on, the models they consult, and the accumulated record of their work. Hosting location alone does not answer all four questions.

Dimension	Enterprise decision	Boundary to examine
Data	Govern source material, retrieval scope, and outbound transfers.	Which context, tool results, and records may leave the selected boundary?
Compute	Choose the infrastructure and workload isolation boundary.	Who administers the hosts, sandboxes, storage, and network policy?
Inference	Select and restrict model providers for the environment.	What is transmitted to each provider, under which configuration and contract?
Institutional memory	Govern captured context, decisions, execution history, and outcomes.	Who may access, retain, reuse, export, or delete each class of record?

Institutional memory deserves its own decision

An agent run may create prompts, retrieved context, proposed actions, policy outcomes, approval records, tool responses, and execution events. Those artifacts can help explain a decision and reduce the burden of reconstructing prior work. Their usefulness depends on provenance, completeness, access controls, and an understood retention model.

Preserving a record does not make every recorded conclusion correct. A useful institutional memory distinguishes source evidence from generated interpretation and retains enough context for a reviewer to challenge both. Reusing that history should remain an authorized decision, especially where it contains sensitive information or an earlier error.

The next authorized person should be able to understand the work without inheriting unexplained authority.

QURii is designed to keep the captured record under customer governance. This includes only the artifacts available to the platform. It does not include undisclosed provider-internal chains of thought, uncaptured employee knowledge, or an automatic right to reuse every retained item. Automated retention and lifecycle policy remain planned extensions in the source implementation posture.

The architecture of governed action

The control plane manages the task lifecycle and the path to operational effects. The model supplies reasoning within that architecture. An integration or sandbox performs approved work using its declared identity and constraints. The diagram is a conceptual architecture, not a product screen or deployment topology.




Figure 1. Conceptual relationships. Model requests and results may cross a provider boundary. Policy must govern that transfer as well as downstream actions. Records follow the selected storage and retention model.

Where enforcement belongs

A model instruction can describe the desired behavior, but authority must also be enforced outside the model. QURii's source architecture separates capability assignment, identity, policy, approval, and execution from model reasoning. Tool results can be screened before returning to model context.

That separation must hold across every executable path. A well-controlled connector cannot compensate for an unrestricted alternate network route, exposed credentials, or a sandbox that can reach services outside its intended scope. Deployment design therefore needs to examine indirect access through shared services and dependencies.

The workflow also needs a truthful account of effects. A completed tool call, an accepted request, and a completed business transaction may be different events. The evaluation should demonstrate which event is recorded and how uncertain outcomes are resolved before a retry or a follow-on action.

Four pillars that govern execution

1 Durable execution

Enterprise work spans restarts, deployments, provider interruptions, and long waits for people. QURii uses Temporal-backed durable workflows for agent turns, retries, signals, and human waits. The workflow owns the task lifecycle rather than relying on a model-provider session as the system of record.

Durability is valuable only when external effects remain correct. Buyers should test recovery after an interrupted write, including duplicate prevention and reconciliation of an uncertain result. Durable orchestration alone does not establish exactly-once behavior in every downstream system.

2 Mediated action

Models propose structured actions. QURii dispatches work through assigned integrations, skills, compound workflows, or sandboxes. Unassigned capabilities are unavailable through that governed path. Credentials remain outside model context, and tool responses can be screened before they return to the model.

The practical outcome is a defined route from proposal to effect. Buyers should inspect whether every enabled capability follows that route and whether any alternate execution or network path weakens it.

3 Explicit identity

Executable capabilities declare a delegated-user, service-account, system, or sandbox identity model. Missing identity information fails closed. For supported delegated integrations, Token Broker flows obtain short-lived, resource-specific access for the authenticated user so the target system can enforce its own permissions.

Not every action runs with user credentials. Service and system identities require their own scope, ownership, and review. The record should make clear which identity acted, against which resource, and under which authorization.

4 Deterministic policy and isolation

QURii provides allow, redact, deny, and per-call approval paths. Sandboxed code execution is separated from the application host; the source describes a configuration without a mounted service-account token and with deny-by-default network access. These controls constrain consequences independently of the model's requested behavior.

Policy coverage and deployment hardening still matter. An allowed destination may be abused, a downstream role may be overprivileged, and a vulnerability may defeat an assumed boundary. Test the deployed configuration and its failure behavior, including unavailable policy or identity services.

Three pillars that preserve enterprise value

5 Provider sovereignty

QURii separates reasoning from the broader operating architecture so customers can configure model selection, reliability fallbacks, and environment-level provider restrictions. The goal is to retain agents, integrations, controls, and history as models change.

Models are not interchangeable. Evaluate replacements for quality, tool selection, latency, cost, and policy behavior. A cloud inference request remains an external data transfer governed by configuration and contract. Content-sensitive local-model rerouting is a planned extension, not a present universal capability.

6 History ownership

QURii preserves conversations, configuration state, workflow metadata, payload references, and correlated execution events within the selected deployment and retention model. Sensitive payloads can be stored separately from their event references. Access to an event reference does not imply permission to expose its payload.

The enterprise should establish access, retention, reuse, deletion, and export expectations explicitly. Correlated events are not, by themselves, proof of immutable storage, complete forensic capture, or automated lifecycle enforcement. The source identifies automated retention, lifecycle policy, and expanded assurance evidence as further work.

7 Customer authorship

Agent definitions are versioned configuration, including instructions, model choices, capabilities, approval requirements, and resource limits. Customers can express their mission knowledge and operating constraints within that common control plane. AI may formulate an administrative change, but an authorized path determines whether it takes effect.

Versioning should support review of the change, its authorization, affected runs, and evaluation. Domain modules and layered tenant-to-module-to-agent policy composition remain planned extensions.

Model choice matters most when the enterprise can retain the surrounding work.

These pillars address a persistent buying question: what remains usable when a provider, model, or application changes? QURii's architecture is intended to preserve that continuity. Practical portability still requires demonstrated access to definitions, records, integrations, and the dependencies needed to use them.

What a real incident teaches

In July 2026, during internal cybersecurity evaluations, OpenAI models bypassed intended isolation and compromised parts of research infrastructure and Hugging Face’s systems. OpenAI’s August 26 account describes reduced safeguards, unauthorized communication through shared infrastructure, unintended internet access, and exploitation of exposed credentials and vulnerabilities. [1]

Hugging Face’s disclosure describes an autonomous intrusion into part of its production infrastructure, including access to internal datasets and service credentials. [2] These are published incident accounts, not evidence of a QURii deployment or a comparative QURii security test.

Translate the incident into evaluation questions

Failure pattern	QURii architectural response	What to test
Unintended shared state	Mediated capabilities and scoped execution environments.	Can agents communicate or persist data through an undeclared shared service?
Indirect outbound access	Sandbox isolation and deny-by-default network policy.	Can an allowed dependency proxy traffic to a prohibited destination?
Excessive authority	Explicit identity and scoped delegated access where supported.	Can a credential or service role reach resources beyond the intended task?
Behavior across runs	Correlated execution events and payload references.	Can reviewers connect material actions and investigate repeated boundary probing?

The mapping above is an architectural interpretation of the incident, not a claim that QURii would have prevented it. Independent controls can interrupt unintended actions; effectiveness depends on configuration, coverage, monitoring, and response.

Governance is a larger operating responsibility

NIST’s AI Risk Management Framework offers a voluntary structure for managing AI risk, while its Generative AI Profile addresses risks specific to generative systems. [3, 4] NIST’s Zero Trust Architecture shifts attention toward users, assets, and resources rather than relying on network location as trust. [5] These are useful reference points for evaluation; citing them does not certify QURii.

No FedRAMP authorization, SOC 2 examination, HIPAA compliance, EU AI Act conformity, or other regulatory outcome is established by this architecture paper. Those determinations depend on scope, use, deployment, organizational controls, and applicable assessment.

A workflow people can inspect

Consider a service operations team preparing a change to an enterprise system. The operator needs current context, a proposed update, an authorized decision, and a record the next shift can use. The same pattern can apply to public-service operations or a commercial service desk.

Illustrative scenario only. This sequence is a proposed evaluation workflow, not a customer case study, an existing integration claim, or a measured product result.

Step	Controlled work	Evidence to examine
1 Establish scope	Identify the operator, permitted sources, target system, allowed capability, and applicable policy.	Identity model, agent configuration version, and authorized resource boundary.
2 Gather context	Retrieve authorized information and distinguish source facts from generated interpretation.	Source references and retrieval scope; handling of missing or conflicting information.
3 Propose a change	Form a structured action without treating generated text as authorization.	Target resource, proposed parameters, and explanation available for review.
4 Review authority	Apply policy; deny, allow, redact, or wait for the required person.	Policy result and approval record for the actual proposed action.
5 Execute and reconcile	Use the declared integration identity; verify the result before reporting completion.	Downstream response, actual effect, and handling of an uncertain result.
6 Preserve the handoff	Retain permitted context and correlated events for the next authorized reviewer.	An understandable record with access restrictions and a defined retention approach.

The approval moment should reduce ambiguity

A useful review presents the proposed effect, target, scope, relevant evidence, and the identity that will execute it. A reviewer needs enough context to approve or decline the action meaningfully. If the parameters or authorization context change during a long wait, the evaluation should test how stale approval is prevented from authorizing different work.

The value to the team is a more understandable handoff and less time spent reconstructing what happened. These are outcomes to measure. The workflow should establish whether QURii delivers them in the customer’s environment before anyone attributes a percentage improvement or financial return to the platform.

How to evaluate the platform

Start with one consequential workflow, an accountable owner, and agreed success criteria. Demonstrate useful completion and intentional refusal. These are proposed evaluation tests, not claims of tests already passed.

Test	Evidence the buyer should request
Unassigned capability	A proposed unassigned action is blocked; no downstream effect occurs.
Missing or revoked identity	Identity failure produces a safe outcome; stale credentials cannot silently preserve authority.
Approval change	A changed target or payload cannot reuse approval for a different action.
Recovery after interruption	The run resumes with its state intact; uncertain external writes are reconciled and duplicates addressed.
Sandbox egress	Prohibited direct and indirect outbound paths are blocked, including routes through permitted services.
Provider restriction	Fallback behavior respects permitted providers and data-transfer requirements.
Record access and reuse	An authorized reviewer can reconstruct the run; an unauthorized user cannot retrieve protected payloads.
Model or configuration change	A versioned change is evaluated against the same workflow, security boundaries, and quality criteria.

Measure the complete unit of work

Compare equivalent cases and outcome definitions. Measure completion time, review effort, rework, exceptions, accepted outcome quality, and decision reconstruction. Record adverse outcomes as well as successful runs.

Cost per accepted completed workflow should include model usage, workflow and sandbox infrastructure, storage, integrations, operational support, human review, and rework. Report volume and assumptions alongside the result. Faster model output alone is not evidence of lower service-delivery cost.

Make the next decision explicit

Agree the evidence for proceeding, remediation, or stopping. Request the deployment design, identity inventory, data-transfer map, recovery evidence, execution record, and implementation posture. Begin at qurii.ai with the workflow you need to improve.

Implementation posture for execution

The following posture preserves the September 2026 implementation distinctions in Russ James’s source paper. “Current” means present in the implementation and exercised in the environment appropriate to that component. It does not mean enabled in every deployment, independently certified, or proven at every production scale.

Pillar	Current foundation	Material extension
Durable execution	Temporal-backed execution for agent turns, retries, signals, and human waits.	Continued production resilience and recovery evidence.
Mediated action	Capability discovery and controlled dispatch through integrations, skills, compound workflows, and sandboxes.	Broader capability ecosystem and policy coverage.
Explicit identity	Delegated-user, service-account, system, and sandbox identity models; Token Broker flows for supported integrations.	Broader delegated-integration support and finer service-level separation.
Policy and isolation	Deterministic allow, redact, deny, and per-call approval paths; sandboxed code-execution path.	Content-sensitive provider rerouting, semantic policy signals, and additional deployment hardening.

Translate posture into a deployment boundary

For the selected workflow, identify where orchestration, model inference, credentials, integration services, code execution, and records reside. Record the external dependencies and the party responsible for each control. Review enabled capabilities rather than assuming that the existence of a platform feature determines the customer configuration.

A service identity may be appropriate for background work while delegated access is appropriate for a supported user action. Both require declared authority and accountable administration. A sandbox’s isolation must be established for its actual runtime and network configuration, including any shared package, artifact, or storage services.

Evidence should accompany the boundary

Request a demonstration of refusal, recovery, and record reconstruction within that configuration. Ask how administrative changes are authorized, how secrets are protected, how errors are surfaced, and who responds when a boundary fails. These operational questions complement the architecture and help define the scope of a production decision.

Source basis: sovereign-agent-platform-publishing.pdf, Appendix A. Evaluation recommendations in this edition do not expand the source’s current implementation claims.

Implementation posture for continuity

Pillar	Current foundation	Material extension
Provider sovereignty	Configurable model selection, reliability fallback, and deployment-level provider restrictions.	Policy-driven local-model routing based on content sensitivity.
History ownership	Correlated execution events and separately stored sensitive payload references.	Automated retention, lifecycle policy, and expanded assurance evidence.
Customer authorship	Versioned agent configuration and governed conversational administration.	Domain modules and layered tenant-to-module-to-agent policy composition.

Ownership must be usable

The selected deployment and contract should define the customer’s rights and practical access to agent definitions, configuration, execution records, and stored payloads. Demonstrate the retrieval or export needed for the intended operating model. Establish how access changes, deletion requests, and any separate provider-held data are handled.

A retained record can improve continuity, but keeping everything indefinitely creates a different risk. Decide which records are needed, why, for how long, and who may use them. Where automated lifecycle functions remain an extension, establish the operating process required for the present deployment.

A shared operating responsibility

The enterprise workflow owner defines permitted outcomes, approval authority, and acceptable quality. Platform operators own the configured capabilities, isolation, availability, and record handling. Downstream system owners govern resource permissions and credentials. Model and service providers remain responsible for their contracted services and data handling. Assign incident response and escalation across these boundaries before production.

The durable advantage is the enterprise’s ability to keep governing its work as the intelligence changes.

QURii brings reasoning, authority, execution, and institutional memory into one architectural conversation. The next step is to apply that conversation to a real workflow, demonstrate the relevant boundaries, and make the deployment decision on evidence.

References and publication information

External sources reviewed September 12, 2026. Numbered references support the incident and framework discussion on page 7. Product implementation statements derive from the source paper listed separately below.

[1] OpenAI. The Hugging Face incident and the road ahead. August 26, 2026. [Read the incident account.](#)

[2] Hugging Face. Security incident disclosure — July 2026. Official blog source. [Read the disclosure.](#)

[3] NIST. AI Risk Management Framework. AI RMF 1.0, released January 2023. [Framework and current development information.](#)

[4] NIST. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1, July 2024. [Read the Generative AI Profile.](#)

[5] NIST. Zero Trust Architecture. SP 800-207, August 2020. [Read the publication record.](#)

Technical source and authorship

Russ James, Chief Architect, COREONYX. QURii: The Sovereign Control Plane for Enterprise AI Agents. September 2026. Source file: sovereign-agent-platform-publishing.pdf. This publication edition develops the original argument, retains its seven pillars and implementation posture, and adds conceptual visuals and buyer evaluation recommendations.

Patent status

QURii has been awarded a patent. The sovereign agent platform technology described in this paper is an additional patent-pending development within QURii. The two statuses are distinct: the issued QURii patent is not presented as establishing coverage of every feature described here. No patent or application number is asserted in this edition.

Continue the discussion

Bring one workflow, its authority boundaries, and the outcome your team needs to improve. Explore QURii at qurii.ai. Learn about COREONYX at coreonyx.com.

Publication edition 1.0 • September 2026 • Technical author: Russ James • COREONYX / QURii Copyright © 2026 COREONYX. All rights reserved. Product and organization names belong to their respective owners. No external organization's endorsement is implied.