

# AI-Native APIs

Designing Platforms for Agentic Autonomy

**Mike Jackson**

Co-Founder & CTO

Varindor Oy

<https://www.varindor.com>

# AI-native APIs: Semantic Middleware, not REST

Executable operating surfaces that agents can discover, use safely, and recover from

Agents need concrete callable actions, not implied workflows



# Agentic autonomy breaks when the API depends on human interpretation

Every implicit API assumption becomes a place where an agent guesses

Humans resolve these questions from external context, agents need answers as data



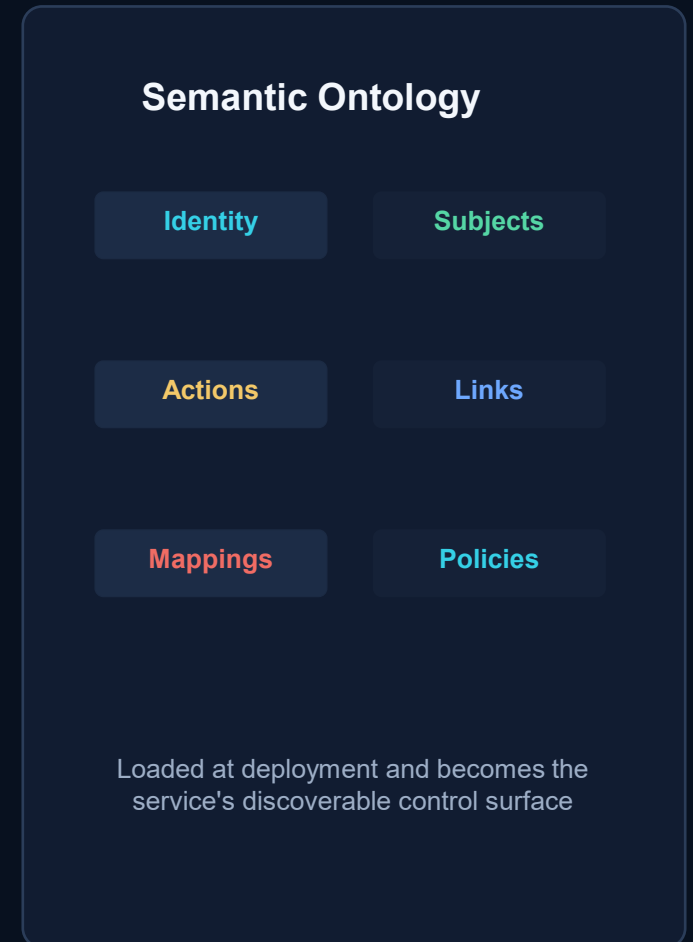
# Move from REST endpoint catalog to Semantic Ontology

The ontology should expose the world the agent acts in, not just methods it can call




**Not documentation. Not SDK folklore. Runtime data.**

Agent code consumes definitions, not assumptions



# Bootstrap the Agents: Runtime Behaviour as Data

An agent should discover identity, role, readable subjects, submittable actions, links, and version drift before it acts



## Same binary

Different identity + role produces a different operating surface

## No hardcoded domain

The role data names the watched subjects and submittable actions

## Hot model refresh

Checksum drift triggers re-bootstrap and atomic model swap

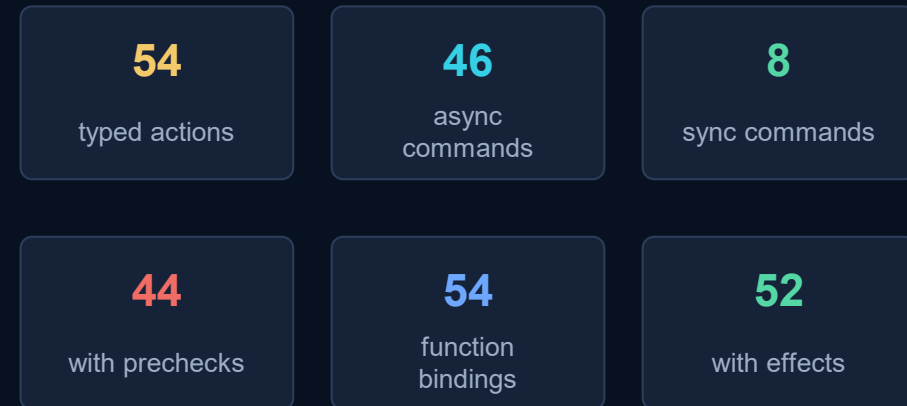
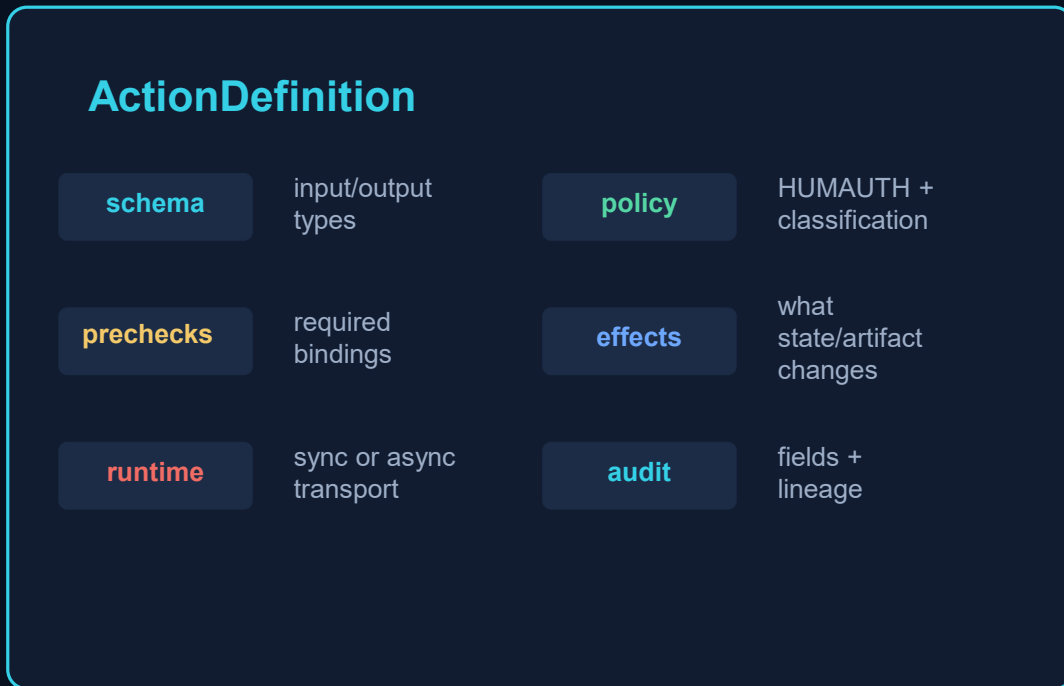
# OODA+ as an Ontology-driven interaction pattern

Observe, Orient, Decide, Act, and Learn are all explicit Ontology actions and data subjects



# An action is a governed command, not a REST endpoint

Schema, policy, prechecks, effects, transport, audit, and lifecycle travel together



The reasoner can propose. The platform still adjudicates.

# Denial as data

A useful refusal tells the agent what failed, where it failed, and what evidence the platform used



failed\_stage

error\_kind

denial\_id

reason

evidence

lineage

**A useful refusal is a machine-readable branch, not a dead end**

# Agents need distributed-systems discipline

Statelessness, idempotency, async lifecycle, drift detection, and topology abstraction matter more than prompt tricks

## Stateless replicas

No session affinity required

## Idempotency

Retries collapse to one decision

## Async lifecycle

Dispatch/result pairs

## Drift checks

Checksum gates re-bootstrap

## Federation

Storage topology hidden


## Audit lineage

Every action leaves a trace

Agent safety is mostly ordinary platform engineering, made explicit

# Expose Ontology Actions and Subjects

Agents need semantically named capabilities, not opaque endpoints



**DISCOVERABLE**

Enumerate the world



**TYPED**

Validate every payload



**SCOPED**

Identity + role bound




**GOVERNED**

Policy outside prompt



**REPLAY-SAFE**

Idempotent retry



**EXPLAINABLE**

Denial with evidence

If an agent has to infer what an API means from documentation, the platform has outsourced control to the model

# Thank You!

<https://www.varindor.com>