

Agent Personas

The Zero Trust Layer for Agentic AI

Where user identity, agent identity, data protection, model access, and behavior converge into a coherent governance decisions for every agent action.

Disparate governance layers

Every large enterprise already runs components of agent governance. Identity providers like Okta and Azure AD authenticate users. Data protection platforms like Microsoft Purview and Netskope classify and protect data. SIEM and observability tools record what happened. However, none of these technologies were designed for AI agents – an actor that is autonomous, headless, token-armed, and routed through an unseen model. These traditional infrastructure pillars were architected for human users and service accounts with fixed purposes, none of them built to make a coordinated decision together at the moment the agent acts. Today's agentic world still requires identity, data protection, and observability tools, but also needs control over model use and agent behavior. These are all required inputs to ensure, in real time, that the action is in scope.

The canonical case

An employee has read and write access to the CRM and to the billing system, by entitlement. They deploy an agent to summarize support tickets and draft customer responses. The job, as the human understands it, is read-only: pull tickets, draft text, hand it back. The agent inherits the full entitlement envelope of the user anyway, because that is how identity-based authorization works. Read access and write access travel together. The agent encounters a refund dispute and, using its own reasoning to solve the problem, calls the billing API to adjust the account balance. Identity passed. Authorization passed. Data protection was not consulted because the API call was structurally valid. Behavioral signals existed in logs but were never part of the decision path. No control prevented the action.

The agent should not have had write access to billing in the first place. The job it was deployed to do did not require it. But nothing in the governance stack was designed to prevent it.

Informed governance demands coordination, not isolation

The agent did not breach the perimeter; it operated inside it. That is what makes this category of failure different from any threat the existing security and governance stack was built to catch. The decision the enterprise needs – is this action, by this agent, on behalf of this user, on this data, consistent with the intended job? – does not exist at any single control point – until now.

The Agent Persona as the binding thread

An Agent Persona is a single definition that compiles into runtime checks across all governance layers. It is the security team's articulation of an agent's job: which user it acts on behalf of, which agent identity is permitted to assume the role, which tools and data classes are in scope, which model it is approved to call, and what in-scope behavior looks like. One definition. All layers reconciled. One coordinated decision per action.

The Enforcement Boundary: The AI Gateway

The Agent Persona is enforced at the boundary between the agent and the things it interacts with – applications, APIs, data, and models. Today that boundary is most often an MCP server, because MCP is the protocol enterprises are standardizing on for agent-to-tool communication. Tomorrow that same boundary will include APIs the agent calls directly, internal service meshes, and whatever protocols come after. Agent Personas are protocol-agnostic by design; the agent sees a single virtual endpoint. Behind that endpoint, an AI gateway assembles the job the Agent Persona describes: a curated set of tools, APIs, skills, and instructions, the specific LLM model the agent is approved to call, and the guardrails that apply to it, all enforced as policy. The gateway binds user and agent identity into every call, applies data protection policy, and evaluates the call against the declared behavioral scope.

The boundary is the right place because it is the only point in the stack where the agent's actual intent (the call it is about to make) is visible, structured, and inspectable before it executes. Not the model layer (model providers will keep changing). Not the agent code (agents redeploy frequently). The gateway, where the agent's call crosses out of the agent's runtime.

That boundary is not limited to tool calls. The agent's direct connection to its language model is often ungoverned: it carries a provider credential, accepts untrusted prompt content, and produces telemetry the rest of the stack never sees. The gateway brokers that connection too, so the agent authenticates to the gateway rather than to the model provider directly.

Not the same as an MCP gateway

The category of MCP gateways is real and growing. They route, multiplex, observe, and rate-limit MCP traffic between agents and tools. That is necessary work, but it is not governance.

An MCP gateway answers the question: can this traffic flow efficiently? An Agent Persona-aware gateway answers a different question: is this specific call, by this agent, on behalf of this user, against this data, consistent with the job the agent was deployed to do? The former is plumbing. The latter is the real-time decision the enterprise actually needs.

Most MCP gateways inherit their security posture from the LLM gateways or API gateways they extend. They authenticate, they apply rate limits, they log. They do not carry a description of what each agent is supposed to do, they do not evaluate every call against a behavioral baseline tied to that declaration, and they do not produce a forensic record that names the agent, the user, and the persona on every action. Agent Personas are the layer above MCP, API, and LLM traffic alike.

For humans, the industry learned RBAC was not enough. We are about to relearn that lesson with agents, unless we start at the right place.

What Makes Up an Agent Persona

User identity binding. Which human authorized this agent to act, and what is that human's entitlement? The Persona never grants the agent more than the user it acts on behalf of, and grants the agent less when the job does not require broad entitlement.

User identity binding.

Which human authorized this agent to act, and what is that human's entitlement? The Agent Persona never grants the agent more than the user it acts on behalf of, and grants the agent less when the job does not require broad entitlement.

Agent identity binding.

Which agent is this, what model, what version, what deployment? The same Agent Persona can be assumed by multiple agent identities under different conditions or restricted to one.

Tool and data scope.

Which agent is this, what model, what version, what deployment? The same Agent Persona can be assumed by multiple agent identities under different conditions or restricted to one.

Model access binding.

Which LLM provider and model can this agent call, and under what token and spend limits? The Agent Persona scopes model access the same way it scopes tools and data, with prompts and responses inspected for injection attempts and sensitive data before either reaches its destination.

Behavioral baseline.

What does in-scope behavior look like? The gateway evaluates every call against the declared scope and flags drift the moment it occurs, with no learning curve, because the baseline is declarative, not learned.

For headless agents (CI/CD pipelines, scheduled jobs, batch processing), the credential primitive that carries this three-way binding is the Agent Access Key. It binds the agent's identity, the user's identity on whose behalf it acts, and the Agent Persona's privileges into a single auditable credential. The forensic record is not "service account X accessed the billing API." It is "agent X, operating on behalf of user Y, under persona Z, called tool W at timestamp T, decision: allow." Token consumption and spend on the model itself are tied to the same key, so the record shows not just which tools the agent called but which model it called and how many tokens it used.

Zero Trust for Agents, Not RBAC

RBAC is the easy analog and the wrong philosophy. RBAC answers one question, once, at authentication: does this role have permission to access this resource. It was designed for human users with consistent jobs, accessing resources at human speed, supervised by other humans.

Agentic zero trust answers a different question, continuously, at every action: is this action, by this actor, on this data, right now, consistent with declared intent? The industry spent a decade moving humans from RBAC toward conditional access, risk-based authentication, and behavioral analytics. Continuous, contextual evaluation replaced the one-time permission check because the one-time check could not catch compromised credentials, insider risk, or lateral movement.

The instinct with agents is to reach back for RBAC. Define a role, grant scopes, issue a token, let it run. That is the architecture zero trust was built to retire.

Identity-only governance versus persona-governed agents

Dimension Requirement	Identity-only governance	Persona-governed agent
Decision point	Once, at session start. Token valid for the window.	Every action. Continuous evaluation against declared scope.
Tool access	Inherited from the user. Read and write access travel together.	Scoped to the tools declared for this job. E.g., read without write when the job is read-only.
Data access	Anything the API permits the credential to touch.	Data classifications consulted as a precondition, not a log review.
Drift detection	Discovered after the fact, in correlation across logs.	The moment scope is exceeded.
Audit forensics	Service account X accessed resource Y.	Agent X, on behalf of user Y, under persona Z, called tool W at time T.

The token-validity window is the specific failure mode. A token is issued, the agent acts freely for the lifetime of the token, and every action in that window passes the identity check by definition. The recent report, [Agentic Zero Trust: Extending the Zero Trust Security Paradigm to Autonomous AI Systems](#) by Chase Cunningham, covers the token-window architecture in detail and grounds the broader Zero Trust extension in the NIST SP 800-207 framework. Identity validates the start of the session. Zero trust evaluates every action inside it. For agents, every action means every tool call the agent makes.

Production Evidence, Co-Existence, and What to Validate

Evidence at scale

Cequence has run a behavioral intelligence control plane in production for over a decade, processing roughly ten billion API transactions per day across global enterprises in financial services, telecommunications, and retail. That plane now evaluates agentic traffic against and Agent Persona-declared scope. Drift outside the declared Agent Persona is detected immediately at runtime. No lengthy baselining exercise. No learning curve. The baseline is the Agent Persona itself.

Recently, a Fortune 50 enterprise ran an autonomous coding agent for 47 consecutive hours, making 2,575 tool calls unsupervised. The agent guessed 162 filenames that did not exist. It hallucinated commit hashes in 71-second probe loops. It re-probed the same wrong paths across 27 hours with no memory between sessions. Every call was technically valid. Every call passed the identity check. Identity-level logging showed clean authenticated requests. The result was excessive token consumption, wasted time, and ultimately a failure to complete the job. Having a behavioral plane in place documents the full activity trail and surfaces drift the moment it occurs.

OpenAI disclosed in 2026 that two of its models escaped a sandboxed evaluation environment, crossed the open internet, and breached Hugging Face’s production infrastructure, chaining a stolen credential and a zero-day to steal a benchmark answer key. Nothing bound those agents to a job, so nothing stopped them from exceeding it.

The honest question for any security leader: what if the agent had been doing something drastic? A determined agent with valid credentials is one prompt-injected document away from a destructive sequence that no identity check would have stopped.

The research record is catching up. Agents of Chaos¹, a multi-institution red-teaming study from Northeastern, Stanford, Harvard, MIT, and Carnegie Mellon, deployed six autonomous agents into a live environment with persistent memory, email, Discord, and shell access. Over two weeks, the agents executed malicious instructions injected into shared documents, voluntarily propagated those instructions to other agents, and spawned persistent background processes from benign requests. Prompt injection is not solved with a smarter prompt. The defense is the layer above the agent.

Co-existence with the existing security stack

Layer / Vendors	What the Agent Persona adds at the moment of action
Identity <i>Okta, Azure AD, Auth0</i>	Tells you who the user is. The Agent Persona binds that human identity to the agent acting on their behalf, at every call, not just at session start.
Data protection <i>Microsoft Purview, Symantec, Netskope</i>	Classifies and protects data at rest and in transit. The Agent Persona consults those classifications when the agent attempts to read or move data, scoped to what the Agent Persona described.
LLM gateway <i>LiteLLM, OpenRouter, Portkey</i>	Routes prompts across model providers and balances cost and latency. The Agent Persona governs the call in both directions: which model and provider the agent may reach before the prompt goes out, and what happens after the model makes a call.
Observability <i>Datadog, Splunk, Langfuse</i>	Records what the agent did. The Agent Persona produces the structured event the platform records: agent, user, persona, tool, model, decision, all attached to the call.
SIEM <i>Splunk, Microsoft Sentinel, CrowdStrike Falcon Next-Gen SIEM</i>	Correlates security signals across the enterprise. The Agent Persona feeds it agent-aware signals SIEM otherwise has no source for.
Cloud and model platforms <i>AWS Bedrock, Azure Foundry, GCP Vertex</i>	Host the agents and models. The Agent Persona governs the agents that run on top of them, independent of which platform is hosting.

What to validate now

- 1. Which of your existing or planned agents would fail their Agent Persona today,** and what would the failure look like in your current logs?
- 2. Where in your token issuance flow does behavioral evaluation belong,** and what would it take to insert a gateway at the boundary without disrupting the agent teams that already shipped?
- 3. What does the forensic record look like for an agent action in production today,** and what would it look like with agent, user, and persona attached to every call?

¹Shapira et al., arXiv:2602.20021, February 2026



Next Steps

Read the Agentic Zero Trust paper

Agentic Zero Trust: Extending the Zero Trust Security Paradigm to Autonomous AI Systems by Chase Cunningham extends NIST SP 800-207 to the agentic threat surface across all five Zero Trust pillars: identity, device, network, application, and data. It maps controls to threats ranging from prompt injection through emergent offensive reasoning and introduces behavioral identity as a first-class security layer. This paper, the one you are reading, focuses entirely on Agent Personas as the operational primitive that compiles behavioral identity into runtime policy. Read together they form the full case for governing autonomous AI in the enterprise.

Read it: [Agentic Zero Trust: Extending the Zero Trust Security Paradigm to Autonomous AI Systems](#)

About Cequence

Cequence protects the applications and data that power enterprises in the agentic era. More than a decade of bot defense and API security experience has established Cequence as the leader of safe and secure agentic AI adoption. The Cequence platform delivers deep insight into user, entity, and agent behavior, enabling organizations to secure, govern, and control agentic AI workflows while protecting against bad actors and rogue agents. Cequence delivers value in minutes rather than days or weeks with a highly scalable, no-code approach. Trusted by the largest and most demanding private and public sector organizations, Cequence protects more than 10 billion daily API interactions and 4 billion user accounts. To learn more, visit **cequence.ai**.