



vs








Willow vs. Runlayer: which control plane fits how your organization runs AI

Willow governs AI for the whole organization. Runlayer secures agent traffic at the MCP layer. Compare both on control, catalog, deployment, and adoption.

Get Started

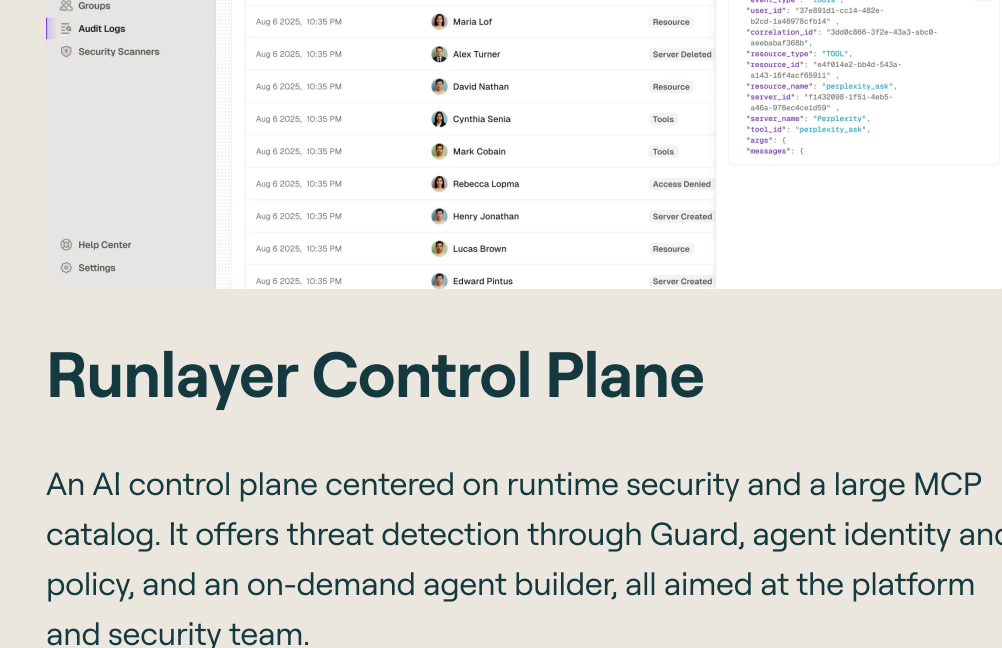
Book a Demo

Trusted by



Willow Admin Console

A managed access and enablement control plane for the whole organization. It provides employee self-service, machine users, fine-grained per-agent and per-action policy, and IT-run identity and lifecycle. It is proven at Wix across roughly 5,000 weekly active users.



Runlayer Control Plane

An AI control plane centered on runtime security and a large MCP catalog. It offers threat detection through Guard, agent identity and policy, and an on-demand agent builder, all aimed at the platform and security team.

The Bottom Line

Archestra is open-source with 659+ MCPs and innovative AI security (Dual LLM, Dynamic Tools). Webrik is a managed platform with employee self-service, machine users, and IT workflows battle-tested at scale.

Willow is the better fit when

You are bringing AI to the whole organization. It gives you employee self-service, a curated and enterprise-vetted catalog, machine users, and endpoint discovery of shadow AI. Its policy is fine-grained and adaptable, and your own IT and security teams can run it themselves. Willow also deploys anywhere, including on-premises, air-gapped, and the EU.

Runlayer is an alternative when

Runlayer is an option if your one priority is runtime threat detection at the MCP layer, the largest raw catalog of servers, and an on-demand agent builder for the platform team, and organization-wide enablement is out of scope.

Pick Willow if

Your goal is to operationalize AI across real teams and tools, with governance your organization can own and adapt, and not only to wire agents to MCP servers.

Willow vs Runlayer, side by side

Dimension	Willow Agent Access Platform	Runlayer
Architecture	Managed access and enablement control plane for the organization	AI control plane centered on runtime security
Deployment options	✓ SaaS, dedicated cloud, on-prem on AWS, GCP, or Azure, air-gapped, hybrid, EU residency	Cloud, self-hosted, and single-tenant through VPC, Terraform, and Helm
Licensing and pricing	Commercial, managed	Commercial, managed
Authentication	OAuth 2.0, OIDC, SAML, JWT, and SSO with Okta, Entra, JumpCloud, Google, and more	OAuth 2.0, OIDC, SAML, JWT, and SSO with Okta, Entra, JumpCloud, Google, and more
Identity and access	RBAC, groups, SCIM, and fine-grained per-agent and per-action policy	RBAC, attribute-based and policy-based access, and runtime-condition policy
Machine users	✓ Generally available service accounts, using an API key or OAuth2 client credentials	Agent Accounts, currently a beta feature that requires access approval
Background agents	✓ Governed autonomous agents, each with its own scoped identity, least-privilege capabilities, and full audit	Runlayer Agents builder with Slack and self-healing; Agent Accounts in beta
Connector and MCP catalog	1,000+ curated, enterprise-vetted connectors with risk scoring	18,000+ MCP servers, continuously scanned
API-to-MCP conversion	✓ Wrap any REST API as a governed MCP	Not a highlighted capability
Employee self-service portal	✓ Connect Panel and marketplace, one-click connection for any employee	Golden-path publishing into the AI clients developers use
IT approval workflows	✓ Approve the catalog once, per-user MCP policy of none, needs-approval, or allow, and governance as code through GitHub	Approve or block through Watch and policy
Shadow AI detection	Endpoint scan agent finds MCPs, skills, and agents on web and local, with risk scoring	Watch finds MCPs, skills, plugins, and clients on managed devices
Vibe app monitoring	✓ Tracks employee use of vibe-coding app builders like Loveable and Base44, and flags when those apps connect to internal systems	Not a highlighted capability
Governed Chrome	✓ Browser extension for visibility into web AI usage and OAuth flows through managed Chrome	Not offered
Observability and audit logs	Full per-call logs, forwarded to Splunk, Coralogix, Loki, CrowdStrike, or a webhook, with configurable retention	Per-request audit, S3 export for SIEM, and spend and ROI reporting
AI security	Built-in runtime and buildtime guards, response PII masking, plus third-party integrations	Guard with AARM Extended Conformance R1 to R9 and patented intent-drift detection
Token optimization	Token and context optimization to cut cost per call	Spend and ROI reporting
Compliance	SOC 2 Type II, GDPR, EU residency	SOC 2 Type II, GDPR, HIPAA
Reference customers	✓ Wix at roughly 5,000 weekly active users, Agora, Riskified, Innovat, Laneweaver	Gusto, Jane, Homebase, Decagon, Opendoor

Get Started

Book a Demo



Who Each Solution Is Best For

Who Willow is built for

Organizations bringing AI to every team without a security-review backlog. It fits when you need:

- ✓ Employee self-service so adoption does not bottleneck on IT
- ✓ Machine users for automation
- ✓ Fine-grained policy that IT and security can run and adapt themselves
- ✓ Shadow-AI visibility before the first incident, including vibe-coded apps like Loveable and Base44, before they connect to internal systems
- ✓ Deployment anywhere your security team requires, including air-gapped and EU

Who Runlayer is built for

Platform and security teams that want an AI control plane centered on runtime threat detection, the largest possible catalog of MCP servers, and an on-demand agent builder. It fits when:

- ✓ Runtime threat detection at the MCP layer is the primary job
- ✓ You want the largest possible catalog of MCP servers
- ✓ You want an on-demand agent builder on the platform
- ✓ The work is securing and scaling agent-to-tool data for a technical team



Governed Background Agents

Autonomous agents with their own identity, scoped to least privilege, and governed like every other call.

Willow approach

- ✓ Every background agent has its own identity and access key, so it authenticates as itself and its capabilities define its blast radius.
- ✓ Scope each agent to the exact tools, skills, and rules it needs, and nothing more.
- ✓ Run agents on any platform, including Custom, Claude Managed Agents, and AWS AgentCore.
- ✓ Every call passes through the gateway under the same guards, policy, and logging as any other traffic.
- ✓ A risk-scored inventory answers which agents can touch a given system, and webhook triggers start an agent when something happens in a connected app.

Runlayer approach

- ✓ Runlayer Agents is an agent builder for creating agents on demand, with Slack integration and self-healing agents.
- ✓ Its agent identities, called Agent Accounts, are currently in beta.

Bottom Line

Both let you run agents, but they optimize for different things. Runlayer optimizes for building agents quickly. Willow optimizes for governing autonomous agents, giving each one its own scoped identity, least-privilege access, and full audit, so security can say yes to unattended work.



Infrastructure as Code

Manage your whole AI configuration the way engineering manages everything else, as code in Git.

Willow approach

- ✓ Sync your toolkits, skills, commands, MCP servers, and clients to a GitHub repository as version-controlled files.
- ✓ Review every change through pull requests, and keep a full history of who changed what.
- ✓ Two-way sync keeps Willow and the repository in step, so a change in either place flows to the other.
- ✓ Authentication secrets are never written to the repository, so the configuration is safe to store.

Runlayer approach

- ✓ Runlayer connects to your identity provider and tools and manages configuration in its own control plane.
- ✓ It does not market a Git-based, pull-request workflow for its full configuration.

Bottom Line

This is where engineering teams lean in. Willow lets platform teams manage AI governance the same way they manage the rest of their stack, in Git with review and history, rather than by clicking through a console.



Self-Service and IT Approval Workflows

Employees self-serve the tools they need, and IT approves once instead of drowning in tickets.

Willow approach

- ✓ Employees browse approved tools in the Connect Panel and connect any AI client in one click, with no IT ticket.
- ✓ IT approves the catalog once, and sets a per-user policy of none, needs-approval, or allow for user-added MCPs.
- ✓ A require-approval guard holds a specific tool call pending a decision, with an approval page that shows what fired and why.
- ✓ When your identity provider deprovisions a user through SCIM, Willow removes that user, and because access flows through their groups, their access to the MCPs and tools they had goes with it.

Runlayer approach

- ✓ Runlayer publishes a governed path into the AI clients teams already use, and admins approve or block through Watch and policy.

Bottom Line

Both keep IT in control. Willow is built so adoption never bottlenecks on IT, because employees self-serve, IT approves the catalog once, and lifecycle changes close themselves.



Guards and DLP

A configurable content layer that redacts, blocks, warns, or requires approval on prompts and actions.

Willow approach

- ✓ Guards inspect every tool call and response at runtime, and inspect content before it is published at build-time.
- ✓ Built-in detectors cover secrets and PII, and you can add regex, JSONata, LLM, or custom function checks.
- ✓ When a guard fires it can redact the sensitive part, warn, require approval, or block, and the most restrictive outcome wins.
- ✓ Build-time guards stop non-compliant content before it ever reaches a user's AI client.

Runlayer approach

- ✓ Runlayer Guard focuses on runtime threat detection, with models tuned for MCP attack vectors and intent-drift detection.

Bottom Line

These solve different halves of the problem. Runlayer leads on runtime threat detection, and Willow's guards are the DLP and policy-as-code layer, configurable to redact, require approval, or block on both on prompts and actions, at build-time and runtime.



ROI and Cost Optimization

See where tokens go and cut them, not just report the spend.

Willow approach

- ✓ Token Usage Analytics shows where tokens come from across MCP servers, toolkits, skills, and tool responses, by team, tool, and use case.
- ✓ Reduce cost by tightening toolkits, moving long skill content into references, and optimizing the tool-response output format, including JSON Compact, CSV, YAML, or TOON.
- ✓ The result is fewer tokens per call, not only a report of what was spent.

Runlayer approach

- ✓ Runlayer's ROI and Observability reports usage, spend, and adoption across users, clients, tools, and agents.

Bottom Line

Both give visibility. Willow goes a step further and actively reduces the tokens each call consumes, so cost control is built into how the platform runs, alongside analytics by team and tool.




Why enterprises pick Willow

We are six to ten months ahead of most companies in AI adoption. More code to production, fewer incidents, real outcomes. Willow is what made it possible to move that fast without slowing down our security posture.



Asaf Yonay
Head of AI Core, Wix



Wix needed a secure, governed way to connect employees and agents to internal tools, documentation, and workflows. With Willow, the AI Core team built the enterprise MCP infrastructure that now supports nearly 600 tools and 300,000+ weekly tool calls across engineering, product, design, HR, finance, legal, and business teams.

60% of enterprise AI rollouts stall on security review
Source: Gartner, 2025

84% / 91% of developers use AI coding tools, with up to 91% of these tools unmanaged or unapproved by IT and security teams
Source: European Cyber Security Organization

Most enterprises have no visibility into shadow AI usage
Source: IDC, 2025