Introducing the ZAgent Framework: The Foundation for Autonomous SASE

June 9, 2026

Zscaler is proud to announce the ZAgent Framework, a new architecture that coordinates a fleet of AI agents across the Zero Trust SASE platform so administrators can automate complex tasks through plain natural language.

The Admin Experience Is About to Look Very Different

The dominant model for enterprise software UI has been stable for decades: you log in, you navigate, you configure, you monitor, you repeat. AI changes the terms of that completely. When an agent can read telemetry, identify an anomaly, trace it to a root cause, and surface a recommended action in seconds, the question stops being "how do I find this in the UI?" and starts being "did the agent already handle it?"

We are already seeing three distinct patterns emerge for how humans and AI systems will interact with security infrastructure:

- First, **conversational:** humans interacting with their security platform the same way they interact with a colleague, through a natural language prompt in whatever tool they are already using, whether that is a dedicated interface, a Slack channel, or a messaging app.
- Second, **generative UI:** when a task is too complex for conversation alone, an agent renders a visualization or workflow on demand, tailored to that specific investigation.
- Third, **fully autonomous**: headless agents connecting directly to APIs, CLIs, and machine-readable tools with no human in the loop at all, handling routine configuration, troubleshooting, policy enforcement, and monitoring in the background.

Most enterprise security platforms today support none of these patterns well. They were built for a console-first world. Zscaler is building for what comes next.

Announcing the ZAgent Framework

The ZAgent Framework is Zscaler's architectural foundation for agentic AI operations across the Zero Trust SASE platform. It is the first step toward fully autonomous SASE, a system that administrators can configure, monitor, troubleshoot, and optimize without logging into a traditional interface.

At launch, administrators interact with ZAgent through a natural language prompt in the Zscaler Experience Center. Make a request in the chat. The ZAgent orchestrator interprets your intent, routes it to the right agent or combination of agents, and returns one coherent answer, regardless of how many systems worked to produce it. Whether a helpdesk person is looking to troubleshoot a performance issue or a network admin is looking to optimize a segmentation policy, the ZAgent knows where to route the request to deliver the desired outputs.

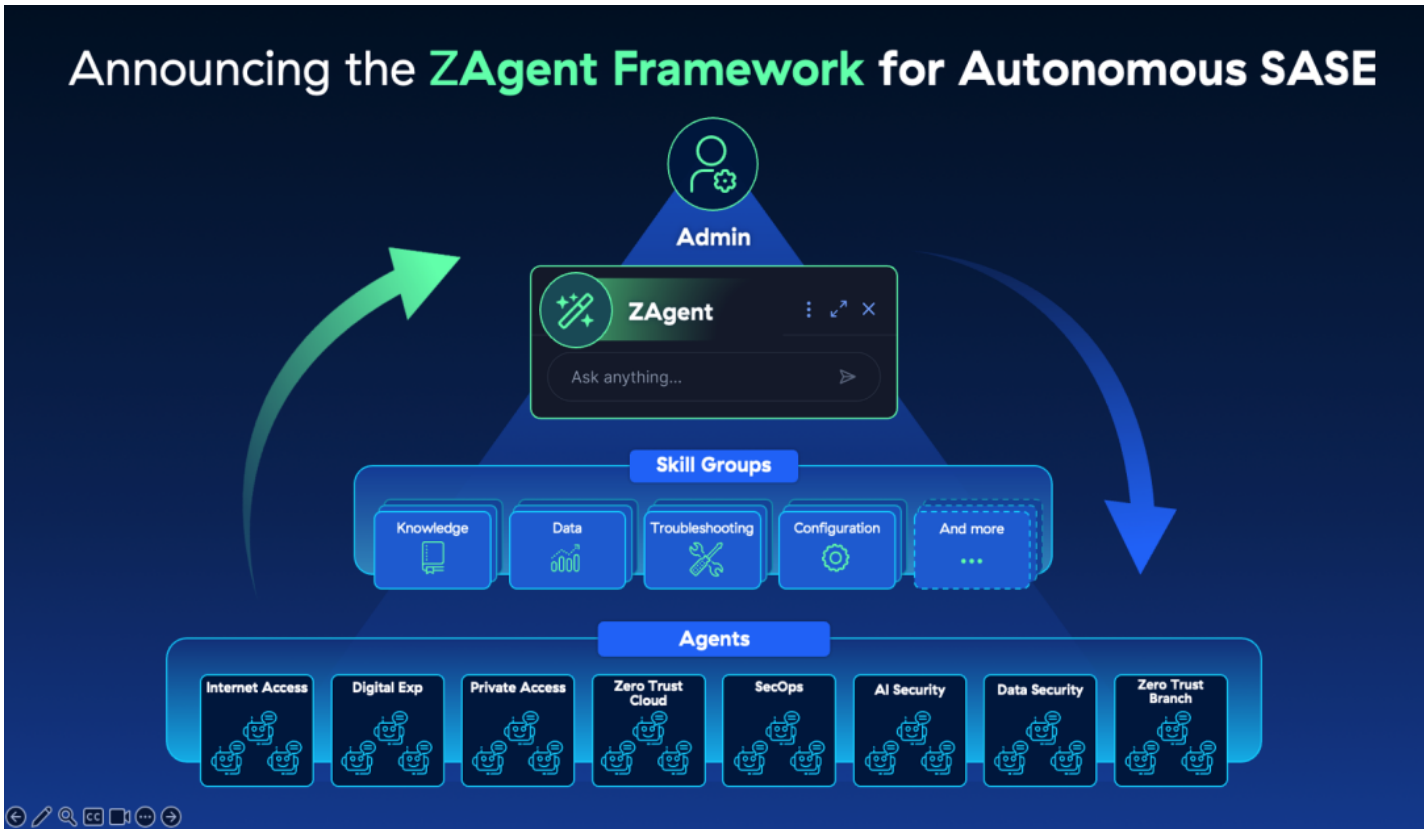

The ZAgent framework delivers multiple benefits to our customers:

1. **Reducing resolution time** for IT and security issues by comprehending the situational context of a challenge and coordinating specialized agents to resolve it.
2. **Simplifying management** of the Zscaler environment by streamlining administration through a conversational interface and automated task execution.
3. **Improving decision accuracy** by correlating context across 500 trillion daily signals and more than 1 trillion AI transactions, reducing false positives and surfacing threats that would otherwise go undetected.
4. **Simplifying audit readiness and risk mitigation** by centralizing governance and guardrails within the Zero Trust Exchange platform.
5. **Expanding team capacity** by automating routine investigation, triage, and remediation workflows so that every administrator operates with the depth and speed of a platform expert.

Specialized Agents, a Shared Set of Skills

The ZAgent Framework is organized around two principles: Agents and Skill Groups.

Agents are domain-specific AI agents, each trained to operate within a defined area of the Zscaler platform. At launch, the framework covers eight areas across our product portfolio:

- Internet Access (ZIA)
- Digital Experience (ZDX)
- Private Access (ZPA)
- Zero Trust Cloud
- SecOps
- AI Security

- Data Security
- Zero Trust Branch

Each agent is a specialist in its domain, with access to the data and tools it needs to act (and no access to anything it doesn't need).

Skill Groups are the capabilities that turn general-purpose AI into Zscaler product experts. Each agent draws on specialized implementations of core skill groups that include (but are not limited to):

- **Knowledge:** Answers questions about products, policies, and configurations.
- **Data:** Surfaces insights from platform telemetry and usage data.
- **Troubleshooting:** Identifies root causes and recommends or executes remediation.
- **Configuration:** Assists with policy setup, segmentation, and platform configuration.
- **Remediation:** Takes action to resolve issues.
- **Workflow:** Orchestrates multi-step operational tasks.

Agents and their skills are activated and orchestrated autonomously by the ZAgent. Administrators don't need to know which agent handled a request or how many collaborated on it. They get the output. And it gets better over time: agents utilize observability to monitor and interpret admin interactions, constantly learning to be able to provide better answers.

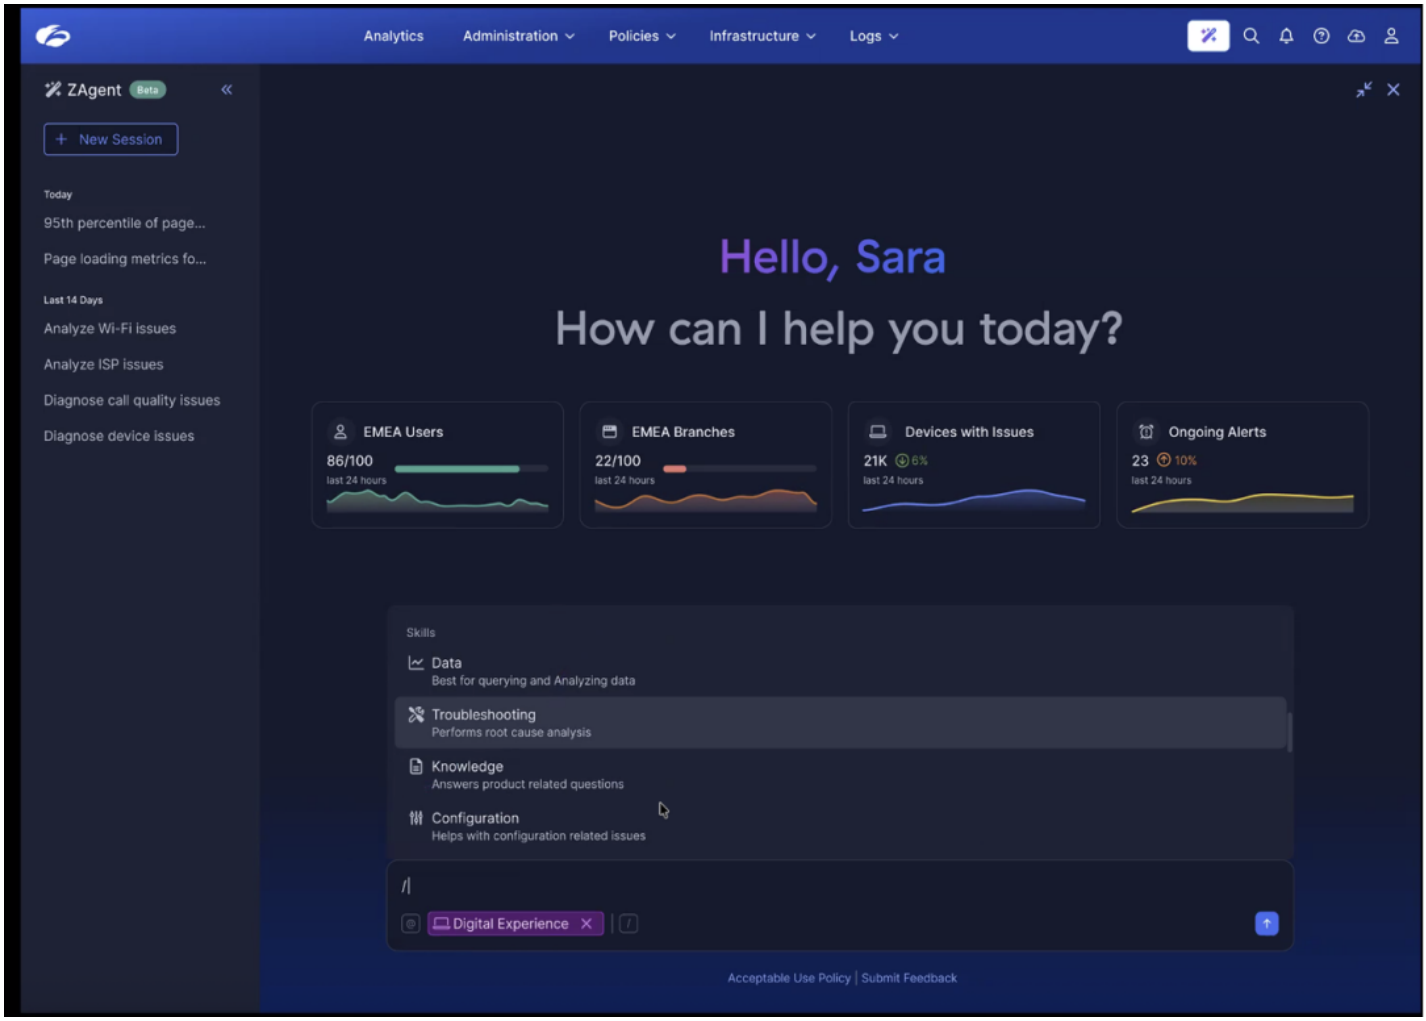

Image

Governance and Compliance

The ZAgent Framework is part of the Zero Trust Exchange platform. The framework has the following controls to ensure we meet local governance and compliance requirements.

1. **Data Residency Controls:** ZAgent Framework data is stored in two geographic regions: United States and European Union (EU), ensuring compliance with regional data sovereignty requirements. Data for every customer's agent follows stringent controls ensuring isolation and no cross-tenant data leakage. Agent requests and responses are processed and stored within the customer's designated region.
2. **Dynamic Role-Based Access Control:** Once a human agent logs in to the admin console, the ZAgent inherits the authenticated human agent's existing permissions. These permissions are evaluated at run time and access can be fine tuned based on existing permissions for the human agent. This ensures the agent is operating strictly with the same boundaries as the human agent it serves, and cannot access resources, policies, configurations etc., beyond what the human agent's role permits at any given time.
3. **Protection from LLM threats:** The agents are protected from threats targeting LLMs, including OWASP Top 10 for LLMs such as Prompt Injections, Training Data Poisoning, and Model Theft. Zscaler is customer zero of AI Guard and Zscaler's broader AI security portfolio to ensure every AI interaction is secure.

AI agentic communication and interaction follows any guardrails and compliance requirements in place within an organization.

ZAgent in Action: Two Early Examples

Our ZAgents already have skills deployed across domains and product areas, and we will be rolling out many more in the coming months. Here are a couple of examples of ZAgent use cases:

ZDX Agent: From Complaint to Root Cause in Seconds

When a user reports a performance issue, the path from complaint to resolution has historically required multiple tools, multiple teams, and significant time. The ZDX Agent's Troubleshooting Skill changes that.

An administrator types: *"Investigate why users in the Northeast are experiencing degraded application performance."* The ZDX Agent builds a multi-step investigation plan, runs it, and returns a summary with findings, supporting evidence, and recommendations in seconds. It rules out endpoints and Wi-Fi and identifies the issue as an ISP problem in the network transit path.

The same agent can investigate an individual user. When a user's performance degrades, the ZDX Agent pinpoints network latency caused by CPU starvation from a video editing process, then brings the administrator into the loop to authorize a remediation job to kill the hung process. The administrator confirms. The action runs. Healthy connectivity is restored.

ZPA Agent: Dynamic Insights from Segmentation Data

Autonomous User-to-App Segmentation, a powerful component of Zscaler Private Access, uses machine learning to identify and fingerprint applications and generate recommendations for app segments and policies. The ZPA Agent extends that capability.

The ZPA Agent lets administrators query segmentation data dynamically, going beyond what is preconfigured in the Experience Center UI. Ask it to show a bar chart of application types across all users and it generates one. Drill into specific application usage by user over time. Tie that output directly to the policies governing access.

Down the line, administrators will be able to use ZAgent to configure and monitor these policies autonomously.

Why Now

Security teams are managing more complexity with the same headcount. The administrators responsible for that problem can't afford to spend time on manual workflows that AI can handle.

The ZAgent Framework addresses that directly. It is a new architectural layer built for the era of agentic AI, sitting beneath the Experience Center today and extensible to any interface in the future.

ZAgent helps ensure the right users have the right access, issues are found and resolved before they become incidents, and your security posture improves without requiring constant manual attention.

What Comes Next

ZAgent will first be accessible through the Zscaler Experience Center. The roadmap extends that same capability beyond the console. Through APIs, CLI workflows, and MCP tools, ZAgent will be able to connect with the AI systems and operational platforms customers already use, including ChatGPT, Claude, ITSM, SIEM, and collaboration tools. You can trigger Zscaler agents from those systems without opening a Zscaler console.

Policy management, insight generation, incident response, configuration at scale: all of it driven by agents working on your behalf.

Learn more about the ZAgent Framework at zscaler.com, or watch the Zenith Live Day 2 keynotes to see the ZDX, ZPA, and SecOps Agents in action.