AI Policy Governance and Enforcement

Build, observe, and enforce strong AI access policies with Kentik

Enterprise leadership is caught in a high-stakes bind: They must enable the massive productivity gains of generative AI while managing the severe infrastructure and regulatory risks that come with unregulated access.

When employees, contractors, or automated agents reach unauthorized AI platforms, they create two distinct operational crises:

The security blind spot

Security teams are blind to "shadow AI" traffic. Because AI traffic is encrypted, legacy network tools cannot identify which endpoints are communicating with risky AI APIs without costly, performance-degrading deep packet inspection.

The compliance liability

Compliance and privacy officers cannot prove that corporate intellectual property or customer personally identifiable information (PII) is safe from being ingested by public LLMs, exposing the company to massive regulatory fines (e.g., GDPR, HIPAA, EU AI Act).

Traditional tools force an unsatisfying fallback. Periodic manual audits are obsolete the moment they are finished, and event-log-based SIEMs are slow, operationally complex, and scale with astronomical data ingestion costs. Security engineers lack the real-time visibility to build precise firewall rules, and compliance officers lack the independent data needed for audit trails.

Solution

Kentik turns network telemetry from next-gen firewalls into a unified, identity-aware record of AI activity across the organization. By correlating real-time flow data with the group, role, and employment-type context already maintained in your identity provider — and tagging traffic against AI services, assistants, and LLM endpoints — Kentik produces a continuously updated view of AI data movement. That view is operationalized for both security engineering and your compliance team through investigation queries, dashboards, alert policies, and natural-language conversations with Kentik AI Advisor, supporting the full policy lifecycle from design through enforcement.

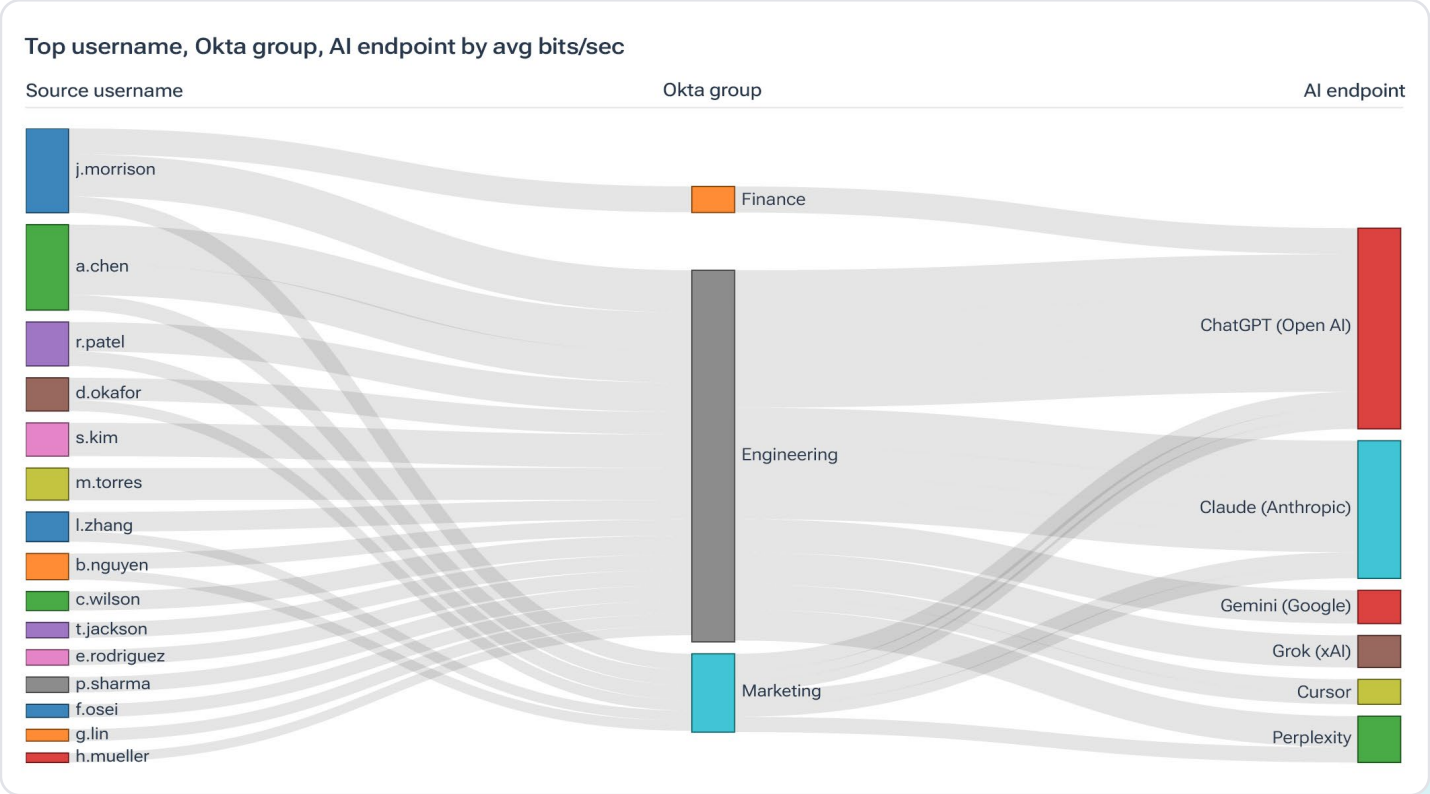

Key use cases

The same data foundation supports the full lifecycle of an AI access policy:

Design smarter AI policy from real user data. Before you can enforce a corporate AI policy, you must understand your current exposure. Kentik surfaces exactly which AI tools are currently processing company data, allowing compliance teams to map shadow AI usage against their regulatory risk frameworks and design informed, defensible data policies.

Validate policy enforcement. Use real-time enriched network flow to validate firewall, SASE, and identity control AI policies are working as intended. Once policy is written and firewall, SASE, and identity controls are configured, prove they are actually holding. Firewall rules drift, exceptions accumulate, and new AI services appear weekly. Continuous network-side visibility confirms what the configuration intends and transforms your network into an automated, continuous verification layer for internal risk assessments and auditors.

Detect shadow AI and policy violations. Alert when a new AI vendor appears in traffic, when a restricted group reaches a sanctioned service for an unsanctioned use, or when contractor activity diverges from policy. And catch bulk exfiltration events to external LLMs before they escalate into public data breaches.

Evolve policy as AI usage matures. New AI services, vendors, and application endpoints appear weekly. Kentik ensures your security policies and compliance reporting evolve at the speed of the market, completely replacing reactive, point-in-time audits.

How Kentik delivers it

The architecture rests on three signals enterprises already collect:

1. Flow data with user identity and end destination identity from next-gen firewalls resolve as usernames, user group, and destination (ie., OpenAI, Anthropic, etc.) context on the flow record.
2. Identity and attribute context from the enterprise identity provider — group membership, role, employment type.
3. AI destination intelligence — application identity, vendor categorization, and curated lists for AI services, assistants, and LLM endpoints.

These are joined inside Kentik as additional dimensions on the flow data. Once configured, every standard Kentik surface — dashboards, queries, alert policies — can group or filter by identity, role, and AI destination across the full retention window.

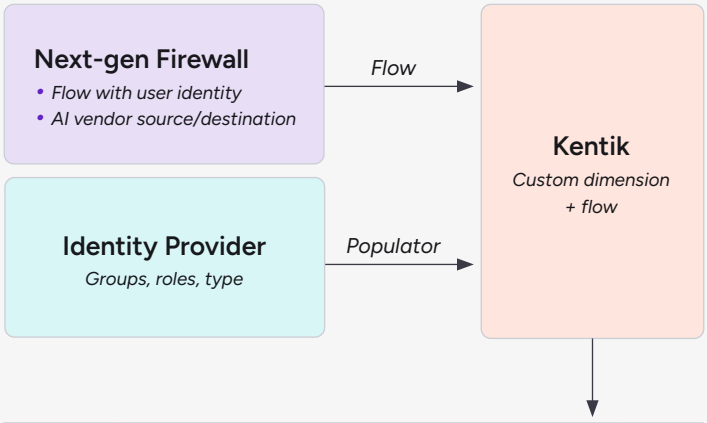

User	Department	Type	AI service	Bytes
jsmith	Marketing	Employee	OpenAI ChatGPT	4.2 MB
rchen	Engineering	Employee	Anthropic Claude	8.7 MB
mlopez	Finance	Contractor	Perplexity	1.1 MB
dkim	HR	Employee	Microsoft Copilot	12.4 MB

Query table in Data Explorer, dashboards, alert policies, and AI Advisor.

Why Kentik

Network truth, not manipulated logs. Flow data is the immutable, real-time metadata record about what is actually happening on the network. Unlike easily bypassed endpoint or application logs that arrive late and require complex parsing, the network is an unalterable source of absolute truth.

Cost structure that fits the use case. Where event-log platforms charge for ingest, retention, and every detection rule, Kentik runs on data customers are already collecting at predictable cost — and supports far more concurrent detections.

Zero infrastructure disruption. Kentik is completely vendor-agnostic. It seamlessly overlays your existing next-gen firewalls, SASE stacks, and identity providers without requiring heavy endpoint agents or complex architectural rebuilds

Conversational by default. AI Advisor is built into the platform, not bolted on. Investigation, reporting, and explanation are part of the product.

One foundation, full lifecycle. The same identity-enriched flow data supports policy design, validation, enforcement, and evolution. The investment compounds as AI governance matures.

Kentik is the network intelligence platform for modern infrastructure teams. Unlike traditional monitoring and observability tools, we demystify complex network operations, enabling organizations to deliver applications and innovation at scale.