

API Security and Management

Alexei Balaganski

July 23, 2025



LEADERSHIP
COMPASS
2025

This Leadership Compass provides a comprehensive overview of the evolving market for API security and management solutions, delivering strategic guidance for identifying products that most effectively meet your organization's requirements. In today's digital landscape, APIs are not only integral to modern applications and services but also serve as critical enablers for AI systems that introduce new vectors for data exfiltration, prompt injection, and misuse. This report identifies leading vendors, products, and innovative offerings that enable consistent governance, security, and compliance across the full API lifecycle.

Contents

Executive Summary	5
Key Findings.....	6
Market Analysis	7
Delivery Models	9
Required Capabilities	10
Integrations	11
Leadership	12
Overall Leadership	12
Product Leadership	14
Innovation Leadership	18
Market Leadership.....	20
Product/Vendor evaluation	22
Spider graphs	22
42Crunch – API Security Platform	24
Akamai – API Security	27
Akto – Akto	30
Axway – Amplify API Management Platform.....	33
Broadcom – Layer7 API Management.....	36
Cequence Security – Unified API Protection.....	39
Check Point – CloudGuard WAF	42
Cloudflare – API Shield	45
Curity – Identity Server	48
Ergon – Airlock Secure Access Hub	51
F5 – Distributed Cloud API Security	54
Forum Systems – Forum Sentry API Security Gateway.....	57

Google – Apigee API Management	60
Gravitee – API Management Platform	63
Imperva – Application Security	66
Kong – API Platform	69
Moesif – Moesif	72
Nevatech – Sentinet API Management.....	75
Pynt – API Security.....	78
Qualys – TotalAppSec.....	81
Salt Security – API Security Platform.....	84
Traefik Labs – Traefik Hub Platform	87
UBIKA – UBIKA WAAP Gateway and UBIKA WAAP Cloud.....	90
Wallarm – Advanced API Security	93
WSO2 – API Manager, Integrator, Identity Server, Asgardeo, Choreo, Bijira, Devant.....	96
Vendors to Watch	99
Akana by Perforce	99
Apache APISIX.....	99
Apiiro	99
Barracuda Networks	100
Boomi	100
Cerbos.....	100
Citrix	101
Data Theorem	101
Fastly.....	101
HAProxy Technologies	102
IBM.....	102
Link11.....	102
Mayhem	103
MuleSoft	103
Ping Identity	103
PlainID.....	104
Radware.....	104
Spherical Defence	104
SmartBear	105

StackHawk	105
Traceable by Harness	105
Tyk	106
Zuplo	106

Executive Summary

Application Programming Interfaces (APIs) have evolved from humble beginnings as technical tools for developers to become the connective tissue of our digital world. They orchestrate business logic, drive automation, and power integrations across mobile apps, cloud platforms, industrial IoT systems, and, in the latest round of disruptive innovation, Generative AI. Every layer now embeds APIs into digital transformation making them essential for building agile, scalable, and intelligent digital enterprises.

The emergence of AI-native applications has accelerated this trend. APIs no longer just support business operations but define them. Modern AI agents interact with each other and external services almost exclusively over APIs. Protecting AI models from prompt injection or misuse is futile if the APIs that interface with those models are left unguarded. In other words, you cannot secure AI without securing APIs: a fact many organizations still overlook in their rush to embrace AI.

KuppingerCole has followed the API space for over a decade, and it has evolved from basic management utilities into a vast and complex ecosystem. As organizations embrace multi-cloud, microservices, and composable architectures, the role of APIs has expanded in both scope and complexity. They are no longer a backend feature, but the product, the risk vector, and often the biggest compliance challenge.




Figure 1: AI growth leading to API complexity explosion

Today's API environments are not limited to REST. They encompass diverse protocols like GraphQL and gRPC, streaming interfaces such as Kafka and MQTT, and asynchronous communication models designed for loosely coupled microservices. These modern API protocols offer performance and flexibility but also amplify the attack surface, increase operational complexity, and expose architectural blind spots.

The recently introduced Model Context Protocol (MCP) is already gaining strong traction. It has emerged as the de facto standard designed to facilitate structured, dynamic, and secure communication between AI agents and external IT systems, including APIs, databases, software tools, and sensor networks. MCP provides a universal interface for context-aware reasoning, real-time decision-making, and modular orchestration of tasks by allowing AI agents to interact with tools and environments in a deterministic, explainable way.

Like REST two decades ago, MCP was not designed with strong security controls in mind, and unchecked implementation of such interfaces can lead to massive data breaches and even more catastrophic scenarios that until recently firmly belonged to the realm of science fiction. Also, emerging technologies like WebAssembly allow applications to execute API logic in isolated sandboxes at the edge or in the browser, introducing new security and observability challenges that traditional tools are ill-equipped to handle.

Even the traditional API gateway, once the centerpiece of API management, is losing its prominence. In fast-moving, containerized, service-mesh-powered environments, centralized gateways are being replaced, or at least augmented, by distributed, autonomous policy enforcement components that operate in real-time and in context. This evolution has fundamentally reshaped the API management and security market. API security has matured into an industry of its own, one that now addresses a broad and growing range of challenges, many of which fall outside the conventional security playbook.

The API security and management market is no longer a niche sector within application infrastructure or network security: it has become a strategic pillar of modern cybersecurity. APIs have emerged as the de facto conduit for exposing business logic, data, and services. They are now both the lifeline of innovation and a critical segment of enterprise security.

Key Findings

- **APIs are the backbone of AI:** Every LLM integration, agentic AI workflow, or autonomous decision system depends on API calls. Most AI-related vulnerabilities, including prompt injection, data exfiltration, or model abuse, are exposed through insecure APIs.
- **Shadow APIs proliferate:** Many organizations struggle with shadow APIs: undocumented, forgotten, or unmonitored interfaces that create unmanaged attack surfaces and compliance blind spots.
- **Business logic attacks are rising:** Unlike traditional vulnerabilities, these attacks exploit the intended functionality of APIs to gain unauthorized access or manipulate data, bypassing rule-based defenses.
- **“Shift Left” meets “Shift Right”:** While API security testing earlier in the development lifecycle (shifting left) remains a growing trend, real-time runtime protection and observability (shifting right) are equally vital. Effective strategies now span the full API lifecycle.
- **Multi-gateway and hybrid API fabrics are now the norm:** Modern architectures often include multiple gateways, service meshes, and other API management tools deployed

across environments, requiring unified management and consistent policy enforcement across them.

- **Edge-native enforcement is gaining momentum:** APIs deployed in edge environments require local policy enforcement and observability, challenging traditional centralized architectures.
- **Compliance is a growing driver:** With regulations like the EU AI Act and GDPR tightening, securing APIs is no longer optional: it is a regulated requirement for any organization handling sensitive personal or operational data.
- **The market continues to consolidate:** In the past year alone, we have seen multiple acquisitions of key API security vendors. These moves reflect a growing trend of integrating API security deeper into broader application delivery and observability platforms.
- **The Overall Leaders in API Security and Management**, in alphabetical order, are 42crunch, Akamai, Axway, Broadcom, Cequence Security, Forum Systems, Google, Gravitee, Imperva, Kong, Qualys, Salt Security, Traefik Labs, Wallarm, and WSO2.

Market Analysis

Traditionally, the API landscape was shaped by vendors offering API gateways, developer portals, and traffic management tools, geared toward performance, availability, and developer enablement. However, as attackers shifted their attention toward APIs as entry points, the industry began to pivot. The rise of API security vendors offering specialized capabilities such as discovery, behavioral analysis, runtime protection, and compliance has redefined the market.

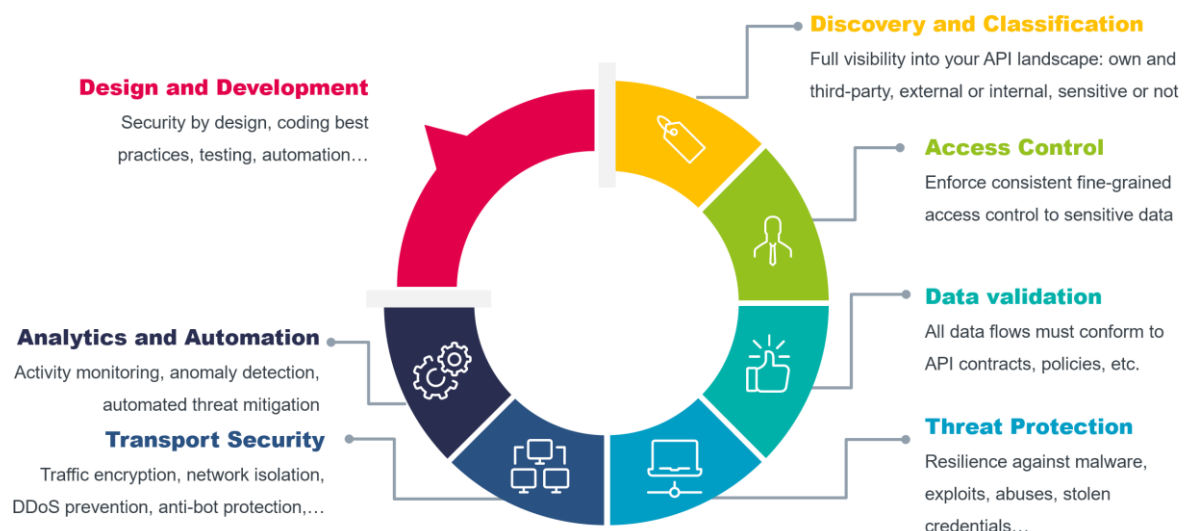


Figure 2: The scope of modern API security

The line between traditional API management and modern API security has blurred to the point of irrelevance. Legacy gateways now evolve into full-featured security platforms, while pureplay security vendors expand their offerings to cover API discovery, posture management, and traffic analysis. API security solutions now operate across every phase of

the API lifecycle, from initial design and automated testing, through discovery and runtime protection, to post-mortem forensics and behavioral analytics. This end-to-end approach reflects the market's shift toward holistic API protection, where security is not a feature, but a strategy.

Several major trends are reshaping the market for the years to come:

The number of APIs in a typical enterprise has grown exponentially. According to Postman's 2024 State of the API report, 92% of organizations increased API usage over the past year, with an average enterprise operates hundreds to thousands of APIs, many of them undocumented or unmonitored "shadow APIs".

The widespread adoption of large language models and AI agents has dramatically expanded the enterprise attack surface. APIs now serve as the communication layer between AI and IT. Risks such as prompt injection, data leakage, and rogue agent behavior must be addressed with multiple cybersecurity controls, and enforcing them at the API layer is the most universal and sustainable approach.

With regulatory frameworks such as GDPR, HIPAA, PCI DSS, and the EU AI Act now enforcing specific requirements for data protection and interface governance (to say nothing about regulations that directly target API or AI management systems like ISO 42001 and NIST SP 800-228), API security is no longer just a best practice, it's a strict requirement for many critical and regulated industries. API logs, access policies, and data flows must be audit-ready, consistently protected, and access to them must be governed by strict and universal policies.

Beyond that, several emerging technologies are beginning to influence the evolution of APIs:

As applications move closer to the edge to meet low-latency, high-resilience requirements, API security must follow. Traditional centralized models cannot scale to protect APIs at the edge. Security enforcement must become decentralized, lightweight, and autonomous. Policy-as-code, local data residency enforcement, and edge-native observability are becoming essential components of modern API security fabrics.

WebAssembly (Wasm) is gaining traction as a secure, portable execution environment for serverless and embedded workloads. APIs delivered via Wasm modules challenge traditional gateway and service mesh architectures. They enable near-native performance, sandboxed execution, and extreme portability, but they also blur the lines between application logic and infrastructure. API security solutions must evolve to inspect, validate, and enforce policies on Wasm-based APIs in real time, even when traditional inspection points are bypassed.

With enterprise IT increasingly accountable for its environmental footprint, API design and infrastructure decisions are coming under scrutiny. Efficient APIs that reduce redundancy and enable smart caching can directly contribute to sustainability goals. Green APIs are becoming a relevant design and operational consideration, especially in large cloud-native environments.

At the same time, the market is increasingly shaped by vendor convergence. Security platforms now integrate with developer tools, CI/CD pipelines, service meshes, IAM, and SIEM systems, forming an “API security fabric” that blends into the broader enterprise IT ecosystem. This convergence is favored by enterprise customers who now expect a platform approach, not a combination of point solutions.

Ultimately, the current market is not just about picking a tool to secure APIs. It’s about building a strategy to manage risk, enforce governance, and maintain compliance across the entire API landscape, including the interfaces that power your AI. In the upcoming years, the market will continue to grow, fueled by increasing complexity. The convergence of application modernization, AI adoption, and regulatory pressure places APIs at the center of enterprise risk and innovation. Vendors that can deliver integrated, intelligent, and adaptable platforms will lead the market.

Delivery Models

API management and security solutions are expected to support highly heterogeneous environments spanning Kubernetes clusters, hybrid clouds, on-premises systems, and edge devices. They are designed to be loosely coupled, flexible, scalable, and environment-agnostic, with a goal to provide consistent functional coverage for all types of APIs and other services. While traditional gateways remain a staple of API management, deployment models now include:

- Service meshes for microservice interconnectivity.
- Inline and out-of-band sensors for runtime traffic analysis.
- SaaS control planes for global policy enforcement and analytics.
- Edge deployments for low-latency, high-resilience enforcement.

These edge deployments increasingly require lightweight, decentralized enforcement mechanisms capable of operating with limited bandwidth, intermittent connectivity, and region-specific compliance constraints.

Pureplay API security solutions that rely on real-time monitoring and analytics may be deployed either in-line, intercepting API traffic, or rely on out-of-band communications with API management platforms. For highly sensitive or regulated environments, customers may opt for fully on-premises deployments. Complex and highly distributed projects, especially those that face data sovereignty requirements, will benefit from edge-based deployments.

However, management consoles, developer portals, analytics platforms, and many other components are usually deployed in the cloud to enable a single pane of glass view across heterogeneous deployments. A growing number of capabilities are now being offered as Software-as-a-Service (SaaS) with consumption-based licensing.

Multi-cloud and hybrid deployment is no longer an option; it is the baseline requirement for modern enterprise security architectures. Vendors must support decentralized enforcement while providing centralized visibility.

Required Capabilities

We are looking for solutions that cover at least several of the following key functional areas, either focusing on more traditional API management or specializing in securing existing APIs (ideally, combining both approaches in a single integrated platform).

API Design and Testing – these functions cover the earliest stages of the API lifecycle such as API contract design, transformation of existing APIs, or modernization of legacy backend services, as well as creating and managing policies that govern API performance, availability, and security.

API Discovery, Classification, and Inventory – without a comprehensive, accurate, and dynamically updated inventory of all APIs across all corporate IT environments (on-premises, cloud-native, hybrid, Kubernetes, etc.) any security program will not be able to provide consistent visibility, governance, and protection across the entire API attack surface.

Microservice Management – traditional API gateways do not scale well for modern distributed architectures and must be augmented with modern service management capabilities such as the Istio service mesh, which provides native connectivity, monitoring, and security that scale for hundreds and thousands of microservices.

Developer and CI/CD Tools – exposing APIs for consumption, providing documentation and collaboration functions, onboarding and managing developers and their apps are among the functions we are looking for here, as well as integrations into existing continuous delivery pipelines of modern application development projects.

Identity and Access Control – Solutions should support multiple identity types, standards, protocols, and tokens while providing flexible dynamic access control capable of making runtime context-based decisions. This does not only apply to the APIs themselves, but management interfaces and developer tools as well.

API Vulnerability Management – discovering existing APIs and analyzing their conformance to API contracts, security best practices, and corporate policies is the only truly proactive approach towards API security. Intelligent prioritization of discovered vulnerabilities by business risk assessment improves both developer productivity and overall security posture.

Analytics and Security Intelligence – continuous visibility and monitoring of all API transactions and administrative activities allow for quick detection of not just external attacks, but infrastructure changes, misconfigurations, insider threats, and other suspicious activities.

Integrity and Threat Protection – securing APIs and services from hacker attacks and other threats requires a multilayered approach to address both transport-level attacks and exploits specific to messaging protocols and data formats.

Strong Internal Security – administrative and developer access to the management console must be secured, with role-based access control implemented across the whole platform and delegated administration capabilities added for scalability and decentralization. Multifactor authentication and audit trail for all activities are recommended.

Hybrid, Multi-cloud Deployment – supporting heterogeneous distributed environments including cloud, containers, microservices, and serverless platforms to be able to provide consistent visibility, analytics, and protection across the entire corporate IT is a critical success factor for any API security solution.

AI Security – expanding the security coverage to protecting application interfaces of large language models against both general-purpose threats like data leaks and LLM-specific risks like prompt injection. Incorporating AI capabilities to improve API security will be considered a major rating boost, too.

Compliance – with the increasing proliferation of API and AI regulations, customers are looking for solutions that can address the burden of complexity of complying with those regulations.

Integrations

Since we do not expect every vendor to be able to provide complete API management and security coverage on their own, seamless interoperability both with the vendor's other own products and with existing third-party products is crucial. A strong focus is put on integration into existing security infrastructures to provide consolidated monitoring, analytics, governance, or compliance across multiple types of information stores and applications.

We expect modern API security solutions to support at least some of the following integrations:

- Popular API management platforms, gateways, service meshes, etc.
- Development tools for integrating into DevOps/DevSecOps processes.
- IAM/IAG platforms for identity management and access governance.
- SIEM/XDR platforms for unified monitoring and security intelligence.
- Additional security solutions like data protection, DLP, WAF, Bot Defense, etc.

Naturally, an API management or security solution also needs to provide its own set of APIs for integration, automation, monitoring, and other purposes.

For more on our methodology and evaluation criteria, refer to the [KuppingerCole Leadership Compass Methodology](#).

Leadership

When selecting a vendor for a product or service, the decision should not be based solely on the information provided in a Leadership Compass. While the Leadership Compass offers a valuable comparison based on standardized criteria and helps identify vendors for further consideration, a thorough selection process requires a detailed analysis and a Proof of Concept (PoC), or pilot phase tailored to the specific needs of the customer.

Based on our research and analysis of the vendor responses, we created various Leadership ratings. The Overall Leadership rating provides a combined view of the ratings for Product, Innovation, and Market Leadership.

Overall Leadership

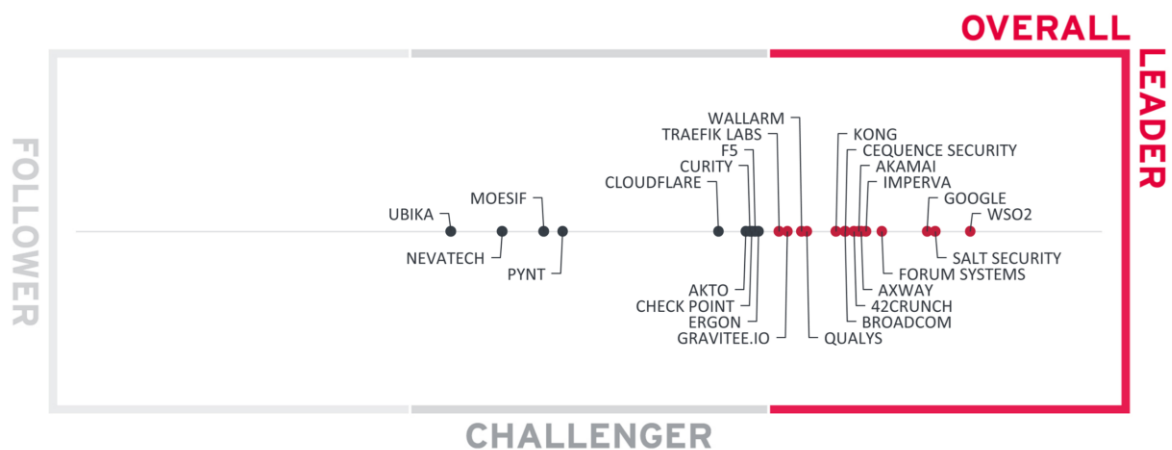


Figure 3: Overall Leadership in the API Security and Management market

The Overall Leadership chart is linear, with Followers appearing on the left side, Challengers in the center, and Leaders on the right. The rating provides a consolidated view of all-around functionality, market presence, and financial security.

However, these vendors may differ significantly from each other in terms of product features, innovation, and market leadership. Therefore, we recommend considering our other leadership categories in the sections covering each vendor and their products to get a comprehensive understanding of the players in this market and which of your use cases they support best.

Among the Overall Leaders, we can find an impressive mix of both established veteran vendors offering enterprise-grade API platforms and smaller, but highly focused and consistently innovative companies – some of which are not even often considered security vendors in the traditional sense.

Akamai, Axway, Broadcom, Google, Imperva, Qualys, and WSO2 are all large corporations with a global presence, strong partner networks, and international customer bases. It is

worth noting that while Akamai was competing against Noname Security last time, it has already finalized its 2024 acquisition, and now, Noname's technology is offered as an integral part of Akamai's API security portfolio. And while Imperva was itself acquired by Thales in late 2023, it still operates as an independent entity in the application security market.

Companies like 42Crunch, Cequence Security, Forum Systems, Salt Security, and Wallarm are representing the pureplay API security solution market, each offering a combination of comprehensive functionality, consistent innovation, and strong market presence enough to be recognized as overall leaders alongside their much larger competitors.

Last but by no means least, vendors like Gravitee, Kong, and Traefik Labs are recognized for their innovative approach, reinventing the traditional API management and making security functions an integral part of the API lifecycle, both from the development and operational perspectives.

Akto, Check Point, Curity, Ergon, F5, and Cloudflare can be found among the Challengers, but they are so close to the upper border of the segment that they have strong chances to achieve Leader recognition next time. In fact, Cloudflare has substantially improved its positioning since the previous version of this Leadership Compass, and the newcomer F5 has only recently entered the API security market.

The rest of Challengers comprises companies, which are either still in the early stage of building strong, international customer bases or are specifically focusing on solving a narrow, but important customer challenge.

While this report was in its final preproduction stage, it was announced that WSO2 has acquired the API observability and monetization Moesif. We have decided to keep Moesif separate in this publication, but in the future, its capabilities will contribute to WSO2's already strong API management portfolio.

There are no Followers in this overall leadership rating.

The Overall Leaders are (in alphabetical order):

- 42crunch
- Akamai
- Axway
- Broadcom
- Cequence Security
- Forum Systems
- Google
- Gravitee
- Imperva
- Kong
- Qualys
- Salt Security
- Traefik Labs
- Wallarm
- WSO2

Product Leadership

The first of the three specific Leadership ratings is about Product leadership. This view is mainly based on the presence and completeness of the required functional capabilities as defined in an earlier section. The vertical axis shows the product strength plotted against the combined/overall strength on the horizontal axis.

The Product Leadership chart is rectangular and divided into thirds. Product Leaders occupy the top section. Challengers are in the center. Followers are in the lower section.



Figure 4: Product Leadership in the API Security and Management market

In this rating, we look specifically for the functional strength of the solutions, regardless of the vendors' current ability to grab a substantial market share. It is worth noting again that,

with the broad spectrum of functionality we expect from a comprehensive API security platform, it is more difficult to achieve a Product Leader status for a smaller company.

This time, we have no less than 21 vendors recognized as product leaders. These, of course, include all the overall leaders described earlier, but also a mix of smaller companies that were unable to reach the overall leader status but are nevertheless capable of offering products with strong, comprehensive functional capabilities. The variety of these solutions highlights the turbulent, rapidly evolving nature of the API security and management market.

42Crunch offers a developer-first API security platform that combines contract-based policy enforcement, CI/CD integration, and runtime protection into a scalable, design-centric solution for securing APIs at every lifecycle stage.

Akamai delivers a comprehensive and deeply integrated API security platform with advanced discovery, behavioral analytics, and hybrid deployment options, powered by its global edge infrastructure and enriched by recent acquisitions.

Akto is a comprehensive API security platform that enables organizations to detect, test, and protect their APIs in real time has been recognized as a leader for its strong automation and runtime protection capabilities.

Axway enables federated API governance across hybrid and multi-vendor environments through its Amplify platform, blending discovery, lifecycle automation, and policy enforcement with strong enterprise and compliance capabilities.

Broadcom offers a mature and extensible API management suite under the Layer7 brand, providing policy-driven lifecycle governance, hybrid deployment, and mission-critical scalability across regulated sectors.

Cequence Security delivers a tightly integrated suite for API protection, bot mitigation, and attack surface discovery without client instrumentation, emphasizing behavioral fingerprinting and rapid deployment.

Check Point integrates advanced AI-powered threat detection and schema enforcement into its CloudGuard WAF platform, enabling proactive API protection with minimal operational overhead in hybrid environments.

Cloudflare offers a comprehensive suite of API security capabilities delivered transparently through its globally distributed edge network, which earned the company a spot among the product leaders for its massive scalability, performance, and integration.

Curity offers a comprehensive and integrated identity and access management platform that enables secure, standards-based authentication and authorization for modern applications and APIs.

Ergon offers a comprehensive API security solution built on its Airlock suite, combining identity-aware access management with advanced threat protection capabilities, recognized for its usability, deployment flexibility, and deep integration of Zero Trust principles.

F5 offers API security services by integrating advanced threat protection, traffic management, and access control across distributed environments, notable for its robust capabilities, scalability, and deep enterprise integrations.

Forum Systems combines protocol mediation, legacy modernization, and GenAI governance in a hardened API gateway architecture, purpose-built for high-assurance environments with zero third-party dependencies.

Google's Apigee provides a mature, AI-enhanced API management platform with deep integration into Google Cloud services, supporting full lifecycle governance, advanced security analytics, and hybrid cloud deployments.

Gravitee delivers a unified platform for synchronous and asynchronous API management with native support for event streams and LLM traffic, underpinned by strong policy governance and hybrid deployment flexibility.

Imperva offers unified API and web application protection with inline enforcement, business logic threat detection, and deep analytics, strengthened by recent integration with Thales and a clear GenAI security roadmap.

Kong combines developer-focused design tools, federated governance, and runtime security in its hybrid API platform, with unique strengths in plugin extensibility and emerging capabilities for LLM and AI agent management.

Qualys extends its industry-leading vulnerability management to APIs through TotalAppSec, delivering automated discovery, testing, and centralized risk scoring across hybrid application environments.

Salt Security enables agentless API visibility and AI-powered threat detection across distributed environments with strong posture governance, CrowdStrike integration, and support for LLM-specific protections.

Traefik Labs offers a lightweight Kubernetes-native platform for runtime API governance with GitOps-driven policy management, observability, and purpose-built features for securing AI-powered APIs.

Wallarm delivers a unified WAAP and API security platform with deep protocol support, inline blocking, and strong ML-based threat detection, particularly well-suited for securing modern, distributed application architectures.

WSO2 provides an open and extensible API management suite combining identity, integration, and analytics, now bolstered by the recent acquisition of Moesif.

The remaining vendors are found in the Challengers segment, mostly as a result of their somewhat narrower focus on solving a specific API-related issue less relevant for traditional API security approaches.

Moesif delivers a sophisticated platform for API analytics, governance, and monetization with strong user-centric monitoring capabilities, but its lack of native API protection features kept it from reaching the leader status in our rating.

Nevatech offers a Windows-native API management platform with strong capabilities in gateway, security policy enforcement, and service virtualization, but its strategic focus on the Windows ecosystem positions it more as a boutique vendor with limited reach.

Pynt delivers a novel approach to API security by integrating automated security testing into CI/CD pipelines, offering developers seamless and actionable insights. However, the company's focus just on shift-left testing substantially limits its scope as a general-purpose API security solution.

UBIKA offers a solid portfolio of API security and management capabilities, particularly in areas like gateway integration and threat protection, but it ultimately landed among the challengers in the Leadership Compass due to weaker developer support, and gaps in visibility and analytics compared to the leaders.

There are no Followers in our product leadership rating.

Product Leaders (in alphabetical order):

- | | |
|---------------------|-----------------|
| • 42Crunch | • Forum Systems |
| • Akamai | • Google |
| • Akto | • Gravitee |
| • Axway | • Imperva |
| • Broadcom | • Kong |
| • Cequence Security | • Qualys |
| • Check Point | • Salt Security |
| • Cloudflare | • Traefik Labs |
| • Curity | • Wallarm |
| • Ergon | • WSO2 |
| • F5 | |

Innovation Leadership

Next, we examine innovation in the marketplace. Innovation is, from our perspective, a key capability in all IT market segments. Customers require innovation to meet evolving and even emerging business requirements. Innovation is not about delivering a constant flow of new releases. Rather, innovative companies take a customer-oriented upgrade approach, delivering customer-requested and other cutting-edge features, while maintaining compatibility with previous versions.

This view is mainly based on the evaluation of innovative features, services, or technical approaches as defined in the Required Capabilities section. The vertical axis shows the degree of innovation plotted against the combined/overall strength on the horizontal axis. The Innovation Leadership Chart is rectangular and divided into thirds. Innovation Leaders occupy the top section. Challengers are in the center. Followers are in the lower section.

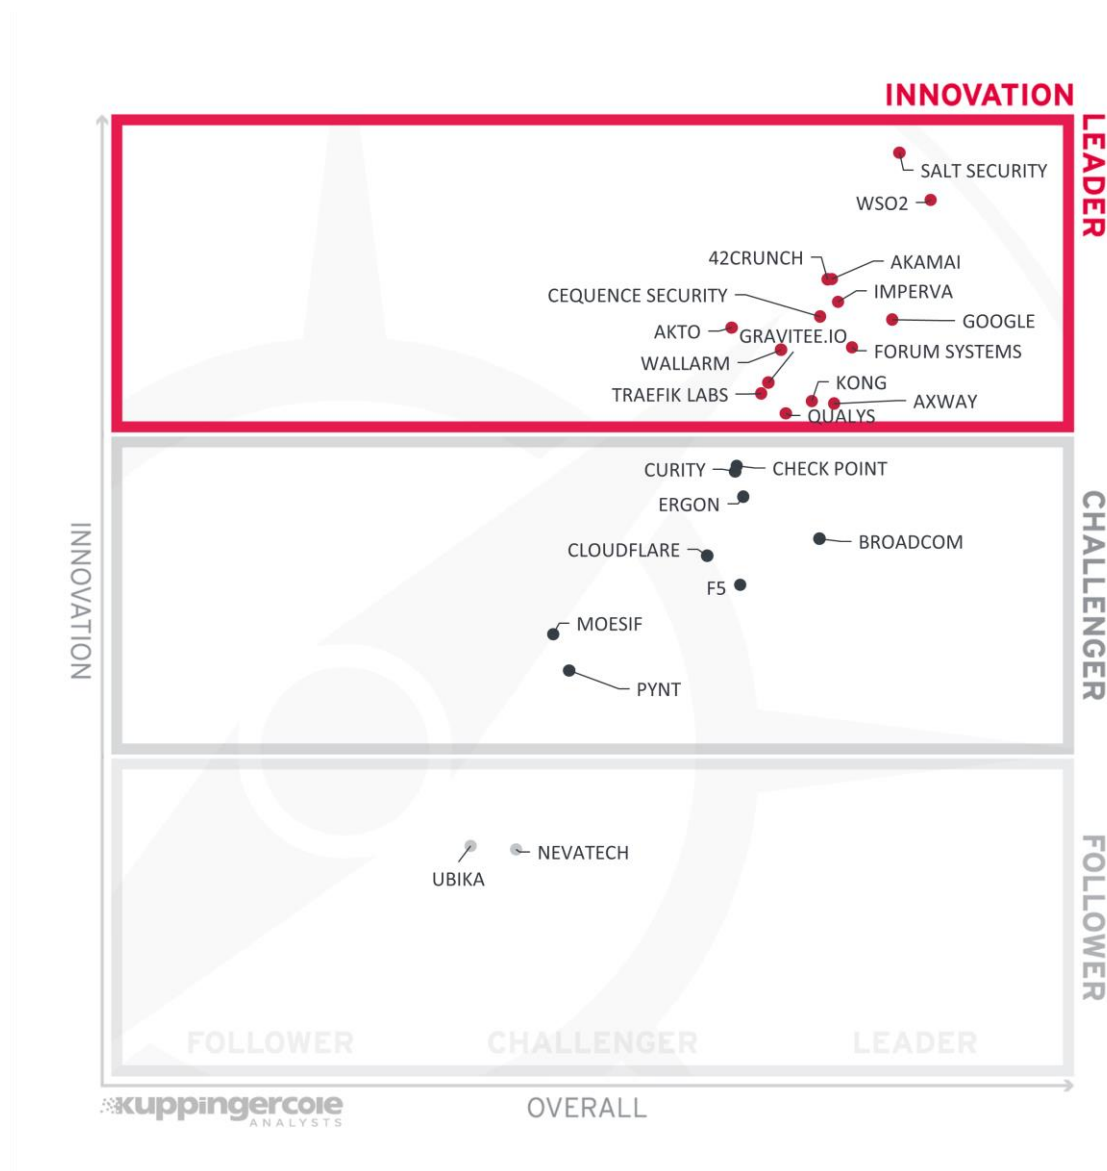


Figure 5: Innovation Leadership in the API Security and Management market

Innovation Leaders are those vendors that deliver cutting-edge products, not only in response to customers' requests, but also because they are driving technical changes in the market by anticipating what will be needed in the months and years ahead. There is a correlation between the Overall, Product, and Innovation Leaders, which demonstrates that leadership requires feature-rich products that are looking over the horizon to bring advancements to help their customers.

Our Innovation Leadership shows a mix of both large and small vendors. This clearly indicates, on the one hand, the massive potential for ongoing innovation in various areas of API security and management, but also that by focusing on a relatively narrow functional area, a small development team can achieve impressive results in delivering useful innovative capabilities in their product. On the other hand, a large R&D investment alone is not enough to deliver a consistent stream of innovation. It is critically important to stay in touch with not just your current customers, but with all the other potential stakeholders, as well as the latest developments in IT, regulated industries, politics, and other areas that influence the future of digital enterprises.

Among the innovation leaders we can again observe many veteran vendors like Akamai, Axway, Forum Systems, Google, Imperva, Qualys, and WSO2, but also smaller and more agile companies like 42Crunch, Akto, Cequence Security, Gravitee, Kong, Salt Security, Traefik Labs, and Wallarm.

A similar mix can be observed among the Challengers, where large companies like Broadcom, Check Point, Cloudflare, and F5 share the space with smaller, more narrowly focused vendors like Curity, Ergon, Moesif, or Pynt.

The remaining vendors can be found in the Followers segment, indicating their somewhat reduced investment into addressing the most recent trends and challenges in API management or security.

Innovation Leaders (in alphabetical order):

- | | |
|---------------------|-----------------|
| • 42Crunch | • Imperva |
| • Akamai | • Kong |
| • Akto | • Qualys |
| • Axway | • Salt Security |
| • Cequence Security | • Traefik Labs |
| • Forum Systems | • Wallarm |
| • Google | • WSO2 |
| • Gravitee | |

Market Leadership

Finally, we analyze Market Leadership. This is an amalgamation of the number and geographic distribution of customers, the size of deployments and services, the breadth and scope of the partner ecosystem, and the financial health of the participating companies. Market Leadership, from our perspective, requires global reach.

In this chart, the vertical axis shows the market strength plotted against the combined/overall strength on the horizontal axis. The Market Leadership Chart is rectangular and divided into thirds. Market Leaders occupy the top section. Challengers are in the center. Followers are in the lower section.



Figure 6: Market Leaders in the API Security and Management market

Again, completely unsurprisingly, we can find large API management vendors like Axway, Broadcom, Google, and WSO2, as well as more security-oriented companies like F5, Imperva, or Qualys among the market leaders – they all have strong financial posture, large partner networks, and customers around the world.

However, since our market leadership ratings depend not just on the vendor's overall financial strength, but their market presence in the API security and management market, some large companies like Cloudflare have not made it into the Leaders segment, simply because API security is still a small part of their entire portfolios.

In addition, we can observe smaller companies like Cequence Security, Forum Systems, or Salt Security among market leaders, simply because of their much stronger focus on just this market segment. Even 42Crunch, a much smaller company, has achieved the Leader status thanks to its strategic partnerships with vendors like Microsoft.

The rest of the vendors populate the Challenger segment, reflecting their ongoing journey towards a larger market position.

There are no Followers in our market leadership rating.

Market Leaders (in alphabetical order):

- 42crunch
- Akamai
- Axway
- Broadcom
- Cequence Security
- F5
- Forum Systems
- Google
- Imperva
- Kong
- Qualys
- Salt Security
- WSO2

Product/Vendor evaluation

This section contains a quick rating for every product/service we've included in this KuppingerCole Leadership Compass document. For many of the products, there are additional KuppingerCole Product Reports and Executive Views available, providing more detailed information.

Spider graphs

In addition to the ratings for our standard categories such as Product Leadership and Innovation Leadership, we add a spider chart for every vendor we rate, looking at specific capabilities for the market segment researched in the respective Leadership Compass.

API Lifecycle Management – here we evaluate the core capabilities of an API management platform, which cover all major stages of an API lifecycle: from architecting an API strategy to developing, deploying, and refining your APIs to daily management and operations, including API monetization.

Deployment and Integration – with the rapid proliferation of API use cases and deployment scenarios, API management platforms must support a wide range of deployment options, from traditional on-premises appliances and static gateways to modern dynamic microservice-based architectures, serverless applications, and IoT, being able to play well together with popular third-party products.

Developer Portal and Tools – exposing APIs for consumption, providing documentation and collaboration functions, onboarding, and managing developers and their apps are among the functions we are looking for here, DevOps and DevSecOps integrations included.

Identity and Access Control – supporting multiple identity types, standards, protocols, and tokens and providing flexible dynamic access control that is capable of making runtime context-based decisions. This does not only apply to the APIs themselves, but management interfaces and developer tools as well.

API Vulnerability Management – discovering existing APIs and analyzing their conformance to API contracts, security best practices, and corporate policies is the only truly proactive approach towards API security. Intelligent prioritization of discovered vulnerabilities by business risk assessment improves both developer productivity and overall security posture.

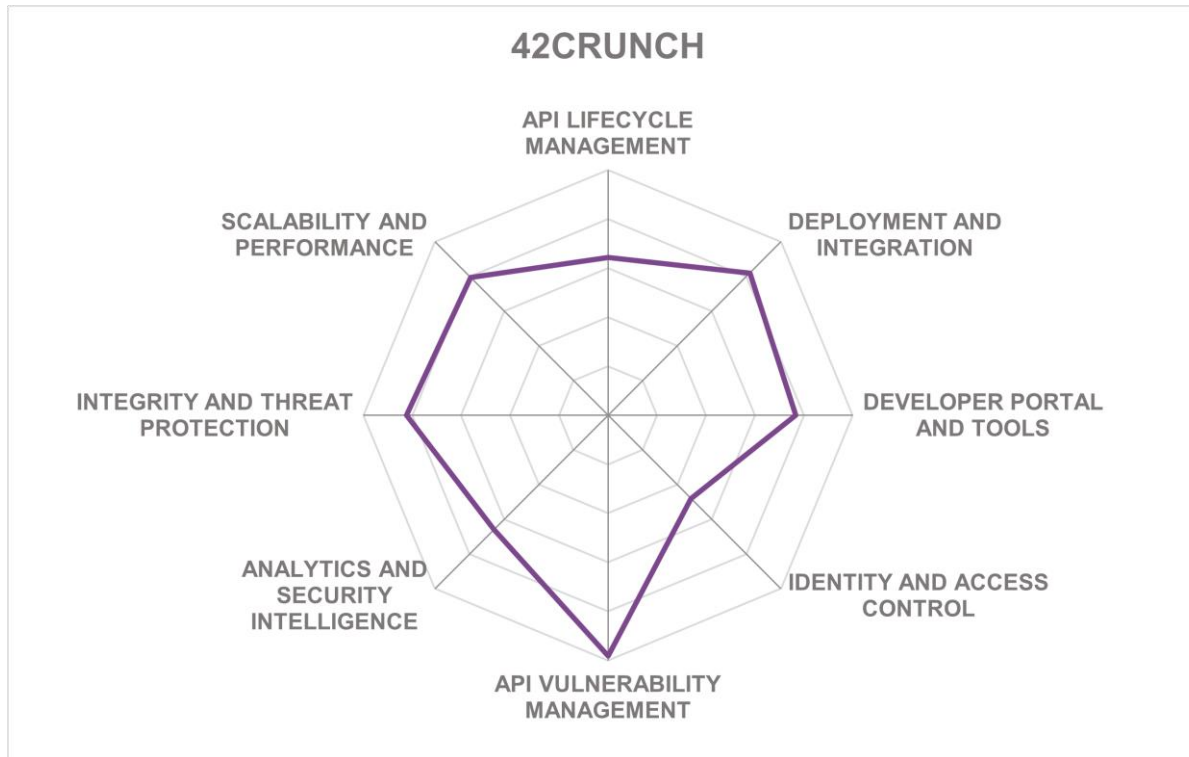
Analytics and Security Intelligence – continuous visibility and monitoring of all API transactions and administrative activities allow for quick detection of not just external attacks, but infrastructure changes, misconfigurations, insider threats, and other suspicious activities.

Integrity and Threat Protection – securing APIs and services from hacker attacks and other threats requires a multilayered approach to address both transport-level attacks and exploits specific to messaging protocols and data formats.

Scalability and Performance – maintaining continuous availability of the enterprise services even under high load or a denial-of-service attack is the most crucial requirement for an API infrastructure. A modern API management solution should also address the challenges of lightweight distributed architectures.

These spider graphs provide comparative information by showing the areas where vendor services are stronger or weaker. Some products may have gaps in certain areas while being strong in other areas. These kinds of solutions might still be a good fit if only specific use cases must be addressed. Other solutions deliver strong capabilities across all areas, thus commonly being a better fit for strategic implementations across complex, heterogeneous IT environments.

42Crunch – API Security Platform



42Crunch is a privately held cybersecurity company founded in 2016 and headquartered in Dublin, Ireland. With a strong presence in the US and Europe and ongoing expansion to Southeast Asia and Australia, the company focuses exclusively on securing APIs through a proactive, developer-centric approach. Its vision centers on empowering organizations to integrate security seamlessly into the API lifecycle, thereby addressing the unique risks APIs pose compared to traditional web applications.

The 42Crunch API Security Platform is designed with a belief that effective protection starts with developers and must be rooted in the structure and semantics of the API itself, typically defined in OpenAPI specifications. The platform supports a shift-left and shield-right paradigm through IDE plugins, CI/CD pipeline integrations that provide continuous feedback and enforcement mechanisms and a runtime protection component. 42Crunch emphasizes security as code and provides comprehensive coverage from API design to runtime enforcement, aiming to make API security a scalable and repeatable practice.

The platform comprises several integrated tools for contract auditing, security testing, and runtime protection. It accepts API specifications in OpenAPI format and automatically performs design-time analysis to identify misconfigurations and risky elements. Developers receive detailed reports and remediation suggestions, including code snippets, directly within their development environments. At the testing phase, the platform offers a suite of scan types (conformance, identity, vulnerability, and drift) that validate APIs against their definitions and detect authorization flaws such as BOLA or BFLA. Security Quality Gates enforce minimum security levels and can be applied dynamically based on tags, roles, or sensitivity levels, helping organizations standardize and scale secure development practices.

For runtime protection, 42Crunch deploys lightweight micro-firewalls that implement a positive security model based on the OpenAPI contract. These enforce fine-grained access control, schema validation, and rate limiting, and can reject malformed or unauthorized traffic with low latency. Additionally, the platform supports GraphQL security auditing and scanning and offers a tool set for automatically generating OpenAPI specifications from API traffic, enhanced by generative AI to produce more readable and accurate documentation.

In addition to core lifecycle capabilities, the platform supports extensive integrations with development, security, and cloud ecosystems. CI/CD compatibility includes plugins for Jenkins, GitHub Actions, GitLab, Azure DevOps, and others. The firewall and scan results can be exported to SIEM tools, while identity-aware testing supports complex credential vaults and role-based access scenarios. A new Azure API Security App enables customers to onboard APIs directly from Azure API Management, streamline setup, and enforce security gates within minutes, reflecting the company's strategic partnership with Microsoft.

The differentiators of 42Crunch lie in its architectural philosophy and practical scalability. The platform uses API contracts as the central source of truth, enabling policy-as-code governance and automated security checks throughout the lifecycle. Its granular security quality gates, extensive IDE and CI/CD integrations, and developer-first UX allow organizations to scale enforcement across thousands of APIs and decentralized teams.

However, the product does not currently offer behavior-based threat detection or anomaly analysis via AI/ML at runtime, focusing instead on a deterministic enforcement model. Some features, such as reporting tailored for regulatory compliance frameworks or delegated administration controls, are also more limited than in some competing platforms.

42Crunch is particularly well-suited for enterprises with a high volume of APIs, complex governance needs, and a mature DevSecOps culture. Its tooling appeals to organizations looking to codify security policies across distributed teams without relying on network-centric detection or post-incident forensics. However, the ease of use of the platform and its freemium pricing model make it an appealing choice for small businesses and individual developers as well. With a clear roadmap focused on developer productivity, automation, and scalable security governance, 42Crunch is a compelling choice for enterprises prioritizing API security-by-design.

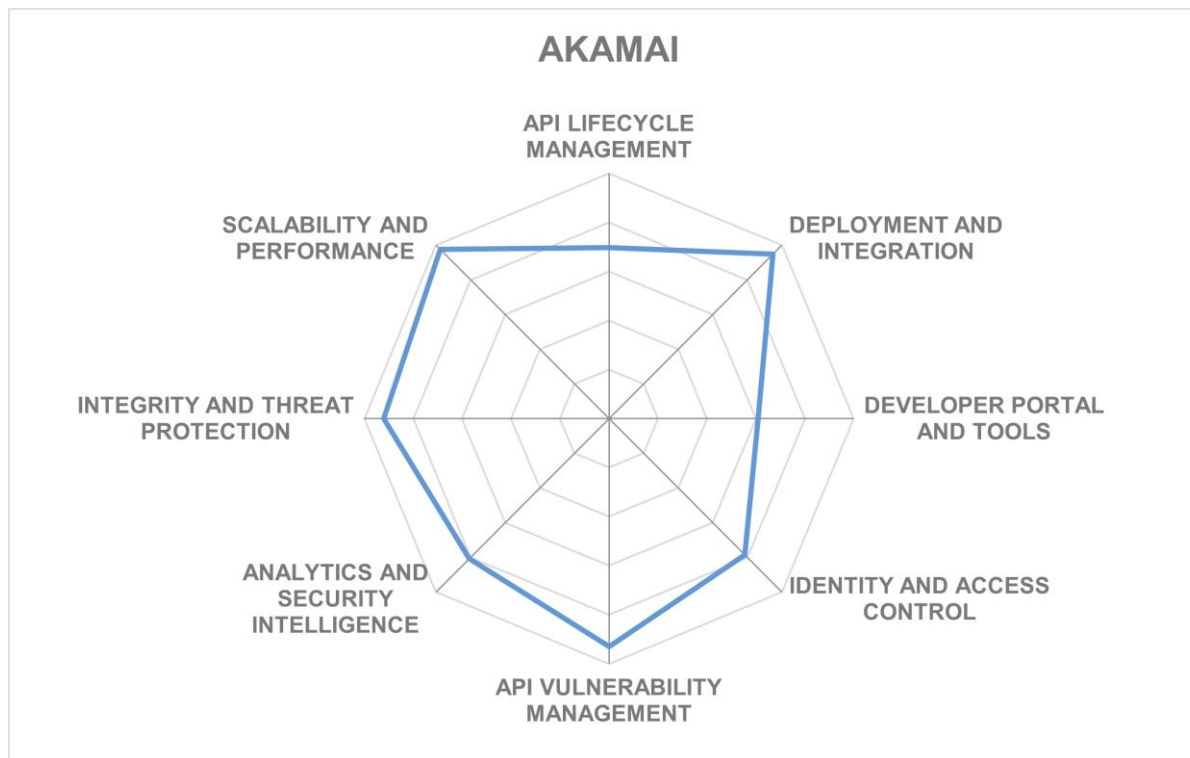
Strengths

- Developer-first approach with native IDE plugins and real-time audit tools.
- Robust CI/CD integration and enforcement of security quality gates.
- Strong coverage of OWASP API Top 10 with design-time and runtime protection.
- Fine-grained, positive security model via OpenAPI-driven micro-firewalls.
- GraphQL schema auditing and scanning support.
- Built-in remediation guidance and AI-assisted OpenAPI generation.
- Flexible deployment models including Kubernetes and hybrid cloud.
- Deep integration with Azure ecosystem including API management.
- Effective enterprise scaling with centralized policy management.

Challenges

- Limited delegated administration and fine-grained RBAC support.
- No support for event-based APIs like Kafka or MQTT.
- AI-based anomaly detection and behavioral analytics are roadmap items.
- Minimal focus on access governance integrations (e.g., SailPoint).

Akamai – API Security



Leader in



Akamai Technologies, founded in 1998 and headquartered in Cambridge, MA, is a well-established provider of content delivery, performance optimization, and security and cloud services. Over its decades-long evolution, the company has transformed from a pioneering CDN operator into a broad cloud and edge platform vendor. With the acquisition of Neosec and, more recently, Noname Security, Akamai has significantly expanded its presence in the API security space, aligning its strategy with the growing demand for integrated, lifecycle-wide protection of APIs that power modern applications.

Akamai's vision in API security is to unify visibility, governance, and protection across cloud-native, on-premises, and hybrid environments while minimizing operational friction. The company positions its API Security offering as a component of a broader platform. This platform includes Web Application Firewall, bot and abuse protection, microsegmentation, and managed detection and response capabilities.

The Akamai API Security platform is a standalone, cloud-native solution, also available in hybrid or fully self-hosted deployments, designed to offer complete API protection from discovery through runtime defense. At its core, the platform collects and correlates information from various sources, including API traffic, specifications, source code repositories, infrastructure policies, and even external scans, to produce a continuously updated API inventory. It supports REST, GraphQL, gRPC, SOAP, and other protocols and is designed to scale across highly distributed environments. Akamai offers tight integrations with its own infrastructure, including native traffic ingestion from its CDN and enforcement integration with AAP. Optional components include a developer-oriented Active Testing module and a managed SOC service for 24/7 incident handling.

The solution provides extensive capabilities across the entire API lifecycle. Discovery and classification of APIs leverage traffic analysis, OpenAPI specifications, source code scanning, and infrastructure metadata, including from AWS and Azure. This allows organizations to uncover shadow and zombie APIs, as well as undocumented or third-party endpoints. The system automatically classifies APIs by sensitivity, data types, and authentication strength, and supports configurable inventory enrichment via labels, tags, and ownership metadata.

The system detects API misconfigurations, outdated specifications, and security weaknesses such as missing or malformed authentication based on traffic analysis. It also offers dynamic security testing of pre-production APIs through its Active Testing engine. Posture management is enhanced by the ability to cross-reference configuration findings with compliance frameworks like PCI DSS, GDPR, and NIST 800-53, supporting both internal governance and external audit requirements.

Anomaly detection is powered by unsupervised machine learning models that baseline normal API usage and detect deviations across authentication methods, input parameters, and traffic patterns. Incidents are scored using the Attacker Confidence Engine, which aggregates historical behavior, IP reputation, geolocation, and threat intelligence into a maliciousness rating to prioritize remediation. Runtime protection can be automated through pre-integrated blocking with Akamai WAF or other security tools.

Akamai's architecture is designed for hybrid and multi-cloud deployments. The system supports agentless ingestion, sidecar proxies, and a lightweight eBPF-based sensor. A unique differentiator is the native connector with Akamai CDN, which allows customers already using Akamai for web delivery to enable API monitoring with a few clicks, achieving full traffic visibility within 20 minutes. Integration breadth is another strength, with more than 300 outbound integrations via the Workato framework, supporting SIEM, SOAR, ticketing, and DevSecOps tools.

Akamai's key differentiators include deep integration with source code and infrastructure to map vulnerabilities to specific repositories and code owners natively and via partners like Apiiro. The platform's support for discovering and classifying LLM and GenAI-related APIs also positions Akamai as a forward-looking vendor for emerging AI workloads. Areas for improvement include expanding capabilities around developer experience and community tooling. The solution is particularly relevant for large enterprises in regulated sectors, which

require scalable, distributed protection for highly diverse API environments. With its flexible deployment model and strong hybrid capabilities, Akamai is also well-suited for customers with complex infrastructure.

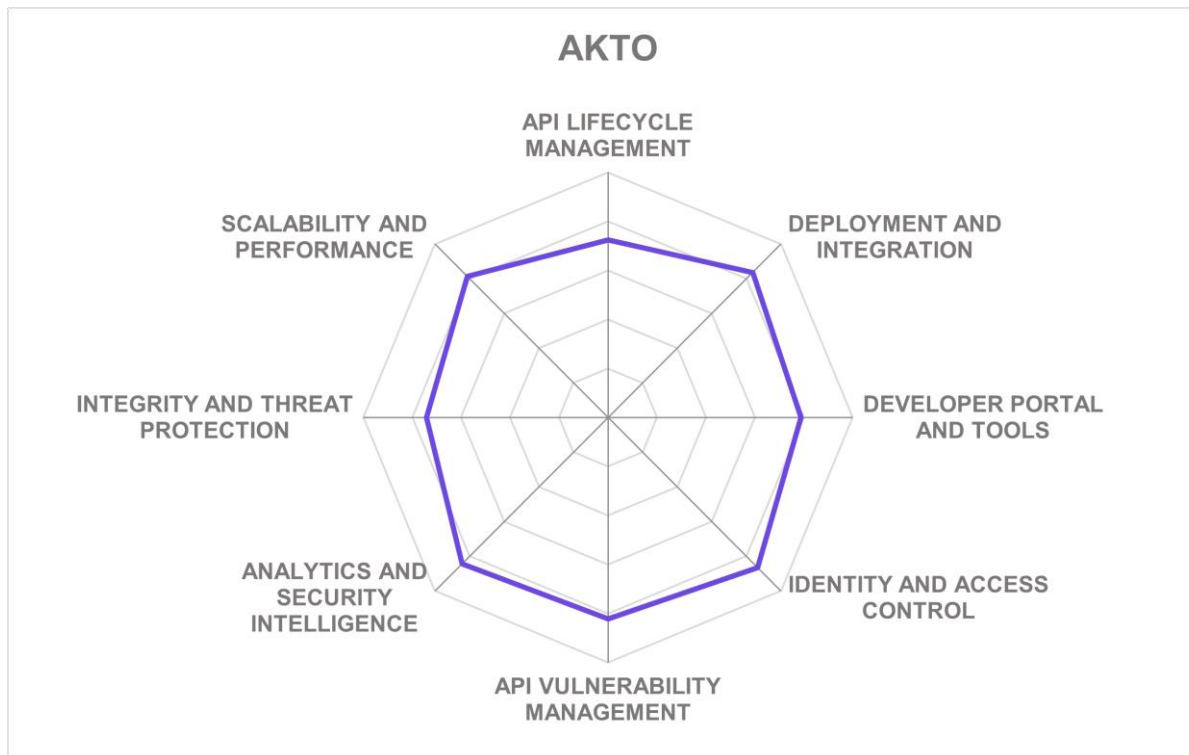
Strengths

- Massive-scale global edge platform enables unprecedented scalability and performance.
- Broad protocol support (REST, GraphQL, gRPC, SOAP, etc.).
- Strong API discovery capabilities combining traffic, specifications, code, and infrastructure.
- Native integration with Akamai CDN for zero-touch traffic ingestion.
- Attacker Confidence Engine for prioritizing incident response.
- Active Testing engine enables shift-left security with pre-production API testing.
- Extensive integration ecosystem with over 300 connectors.
- Flexible deployment options: SaaS, hybrid, on-premises, or hardened appliance.
- Managed SOC service for API-specific threat analysis and response.

Challenges

- No dedicated developer portal for key provisioning or API management.
- Limited built-in support for event-based protocols (like Kafka or MQTT).
- Does not currently support protocol transformation or legacy SOA repackaging.
- Brand awareness is still adjusting after the Noname Security and Neosec acquisitions.

Akto – Akto



Leader in



Founded in 2022 and headquartered in San Francisco, CA, Akto is a relatively young but ambitious vendor in the API security market. Positioned as a pure-play API security company, Akto aims to support modern AppSec teams by enabling strong visibility and security throughout the API development lifecycle. The company's go-to-market strategy emphasizes proactive security across all stages of DevSecOps and targets industries with large and complex API infrastructures.

Akto's strategic focus is to build a unified, lifecycle-spanning API security platform that combines visibility, posture management, vulnerability assessment, and runtime threat detection. Rather than concentrating solely on reactive defenses at runtime, Akto deliberately started with API discovery and security testing as its foundational capabilities. Over time, the product has evolved to offer broader coverage, from identifying APIs at the source code level to monitoring runtime traffic in production. The company positions itself as

the only vendor to provide “code to runtime” API security, and while this claim should be considered with nuance, it reflects Akto’s intent to tightly integrate security into DevSecOps pipelines.

The Akto Platform is delivered as a single, integrated solution available both as a SaaS service and for self-hosted deployment, supporting hybrid use cases where data processing happens within customer environments while dashboards are managed in the cloud. Akto combines API discovery, vulnerability scanning, posture management, and threat detection. The platform includes over 1,000 out-of-the-box security tests tailored specifically for APIs and a broad range of injection and access control vulnerabilities. These tests can be executed continuously, on a schedule, or in CI/CD pipelines, and they support detailed remediation workflows.

API discovery engine leverages over 50 traffic and code connectors to uncover APIs from source code, gateways, Kubernetes clusters, and cloud services. Akto can build a comprehensive inventory of internal, external, partner, and third-party APIs, detecting zombie and shadow APIs as well as identify sensitive data exposure patterns. Its ability to extract API information directly from source code repositories like GitHub or Bitbucket significantly broadens its coverage of pre-production APIs.

Once discovered, APIs are tested using Akto’s security test library, which includes capabilities for dynamic analysis and regression testing. The platform supports tokenized and contextual access control validation using roles, credentials, and login steps that can be customized to match real-world authentication scenarios. Test results are presented with contextual evidence, including reproduction steps and recommended code-level fixes, streamlining remediation for development teams.

Akto further extends its coverage into runtime protection by integrating with traffic sources and monitoring API calls for anomalous behavior. It uses machine learning and rule-based techniques to identify threats such as Broken Object Level Authorization (BOLA), Server-Side Request Forgery (SSRF), and mass assignment attacks. The platform can forward threat indicators to WAFs and SIEM systems, although it does not perform inline blocking itself. This approach avoids the risks associated with in-path deployments while still allowing for automated threat response through integrations with third-party defenses.

To support compliance and governance, Akto provides centralized dashboards and risk scoring, with built-in reporting aligned to major regulations such as GDPR, HIPAA, SOX, and PCI DSS. The platform offers visibility into authentication types, API types (including REST, SOAP, GraphQL, and gRPC), and data sensitivity. It includes a management UI for viewing API security posture over time and integrates with ticketing systems like Jira for incident tracking.

A unique differentiator for Akto is its capability to analyze API definitions directly from source code repositories using AI-assisted logic. This enables discovery and risk evaluation of undocumented or legacy APIs that may never reach production, helping organizations to deprecate unused endpoints early. The platform’s recent introduction of agentic AI modules promises to enhance various workflows, such as vulnerability triage and test generation.

Akto primarily serves midsize and enterprise customers, especially in regulated sectors such as banking, insurance, and healthcare, where understanding and controlling API risk is a top priority. The solution's coverage of both pre-production and runtime environments makes it particularly relevant for teams adopting a DevSecOps methodology. Its lightweight, agentless deployment model and support for hybrid environments also make it a suitable option for organizations with complex infrastructure across cloud and on-premises.

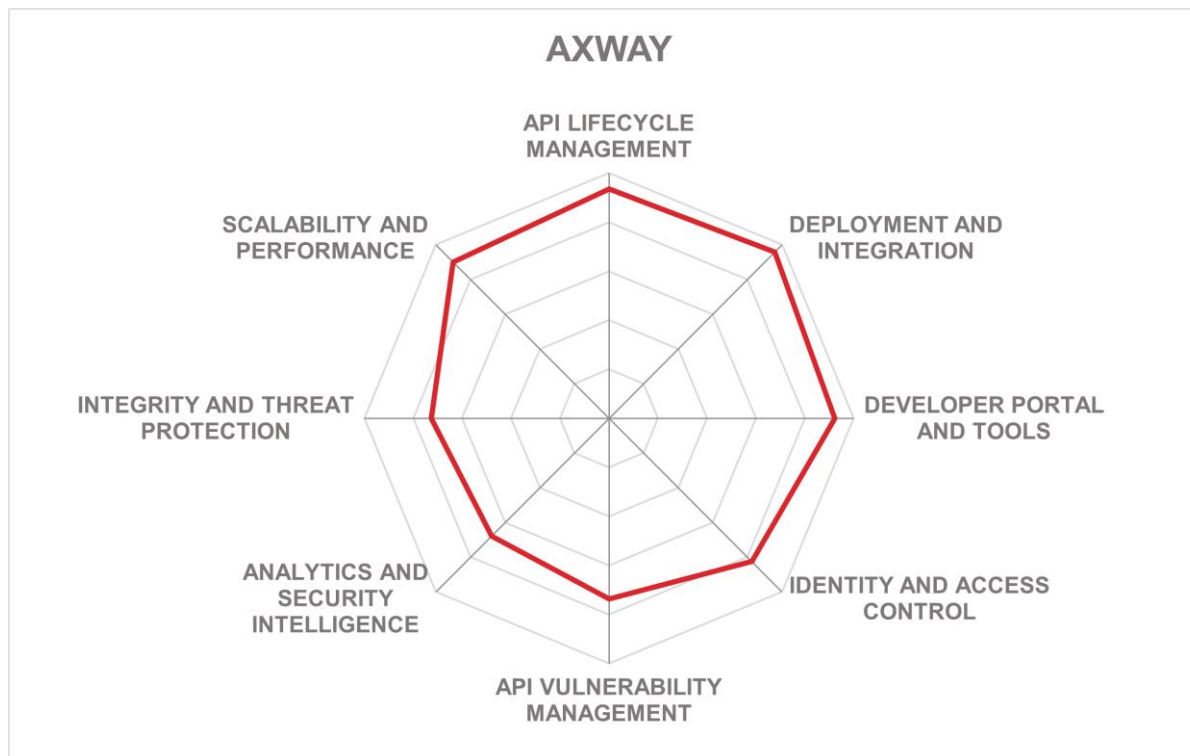
Strengths

- Unified platform covering API discovery, testing, posture management, and runtime threat detection.
- Extensive out-of-the-box test library with over 1,000 API-specific security tests.
- Unique API discovery from source code using language-agnostic analysis.
- Comprehensive support for API types (REST, SOAP, GraphQL, gRPC, etc.).
- Deep integration into DevSecOps pipelines and CI/CD workflows.
- Dynamic testing capabilities with contextualized remediation suggestions.
- ML and AI capabilities for anomaly detection, test generation, and vulnerability prioritization.
- API risk scoring and dashboards tailored for AppSec teams.

Challenges

- Market visibility and ecosystem reach are still limited compared to established competitors.
- Lack of inline blocking capabilities may limit protection in high-security environments.
- Dependency on external WAFs for active mitigation may complicate deployments.
- Discovery of third-party APIs relies primarily on traffic and is limited in depth.

Axway – Amplify API Management Platform



Leader in



Axway, founded in 2001 and headquartered in Scottsdale, AZ and Paris, France, is a French-American software company offering digital transformation solutions including API management, managed file transfer, and B2B integration. After the recent acquisition of Sopra Banking Software, Axway now operates as an independent part of the 74Software group. The company maintains a global presence with key offices in Europe, North and South America, and Asia-Pacific. Over the years, Axway has expanded its product portfolio both organically and through acquisitions, including the notable addition of DXchange.io, which brought advanced integration capabilities into the fold.

Axway's strategic vision in the API security and management space is centered around addressing the fragmentation and complexity enterprises face in modern digital ecosystems. The company emphasizes a federated and unified approach to managing APIs across multiple gateways, cloud platforms, and integration styles. This vision is supported by three

core components of the Amplify platform: Amplify Engage, Amplify Fusion, and the Amplify Gateway. Together, they bring consistency, visibility, and control across the full API lifecycle regardless of whether APIs reside in legacy systems, modern clouds, or third-party environments.

The Amplify API Management Platform provides an integrated set of capabilities spanning API gateway services, a universal developer portal, lifecycle automation tools, and a low-code/no-code integration layer. Amplify Engage serves as the platform's centralized portal for API productization, discovery, and consumption, supporting APIs from Axway and third-party gateways like AWS, Azure, Kong, MuleSoft, and others. Amplify Fusion provides orchestration, data mapping, and composite API creation capabilities, leveraging iPaaS strengths and an embedded Envoy-based gateway. The traditional Axway API Gateway remains at the core of the enforcement layer, enabling fine-grained security, traffic shaping, and compliance enforcement across various API protocols.

The platform offers comprehensive support for API design, transformation, and policy creation. Integration with tools like Stoplight supports collaborative contract-first API development, while Amplify Fusion allows organizations to build and expose composite APIs without code. The platform's policy engine includes over 200 prebuilt operators and graphical configuration tools, facilitating quick enforcement of authentication, encryption, quota management, and compliance rules.

Amplify's discovery and classification capabilities are powered by a rich agent framework, enabling visibility into both managed and unmanaged APIs across diverse environments. Discovery agents connect with source code repositories, CI/CD pipelines, API gateways, and observability tools, consolidating inventory into a unified service registry. Axway integrates with runtime discovery platforms such as Traceable to identify APIs in production traffic that may not be registered in governance systems. This federated approach ensures that even rogue or shadow APIs can be discovered, validated, and remediated proactively.

Axway supports comprehensive threat mitigation, including embedded Web Application Firewall (WAF) functionality, schema validation, content filtering, rate limiting, and access management integration with major IAM platforms. The platform addresses all OWASP API Security Top 10 threats and includes fine-grained access controls supporting all major standards. Fusion further enables integration flows that incorporate API security enforcement as part of orchestration pipelines, closing the loop between integration and protection.

Analytics and observability capabilities are integrated across all modules, with Amplify providing real-time dashboards for usage metrics, API health, consumer engagement, and infrastructure monitoring. Custom reports and alerts can be created through a visual interface, and integrations with SIEM tools are available. Amplify Engage additionally supports governance workflows, SLA enforcement, and subscription management for internal and external API consumers.

Axway supports hybrid and multi-cloud deployment models, offering its solution as SaaS, self-managed, or through private cloud deployments. The platform's modular architecture

enables customers to deploy only the components they require. A modernized user interface, Amplify Studio, further unifies management across control planes and data planes, supporting both traditional and cloud-native topologies.

Although the platform currently does not incorporate ML-based threat detection natively, it partners with external providers to offer these capabilities. On the generative AI front, Axway is investing in capabilities for API documentation generation, integration flow design, and LLM brokering. The platform could benefit from expanding its native AI-driven threat detection capabilities beyond partner integrations. Anomaly detection, risk scoring, and autonomous response still depends on external tooling.

Axway's strengths lie in its federated governance model, ability to unify API inventories across diverse technologies, and strong support for the entire API lifecycle. The platform is especially appealing for enterprises managing large-scale API ecosystems across multiple vendors or integration teams. Use cases such as internal API marketplaces, composite API orchestration, or Open Banking compliance are particularly well-supported. It is well-suited for regulated industries and complex enterprise environments.

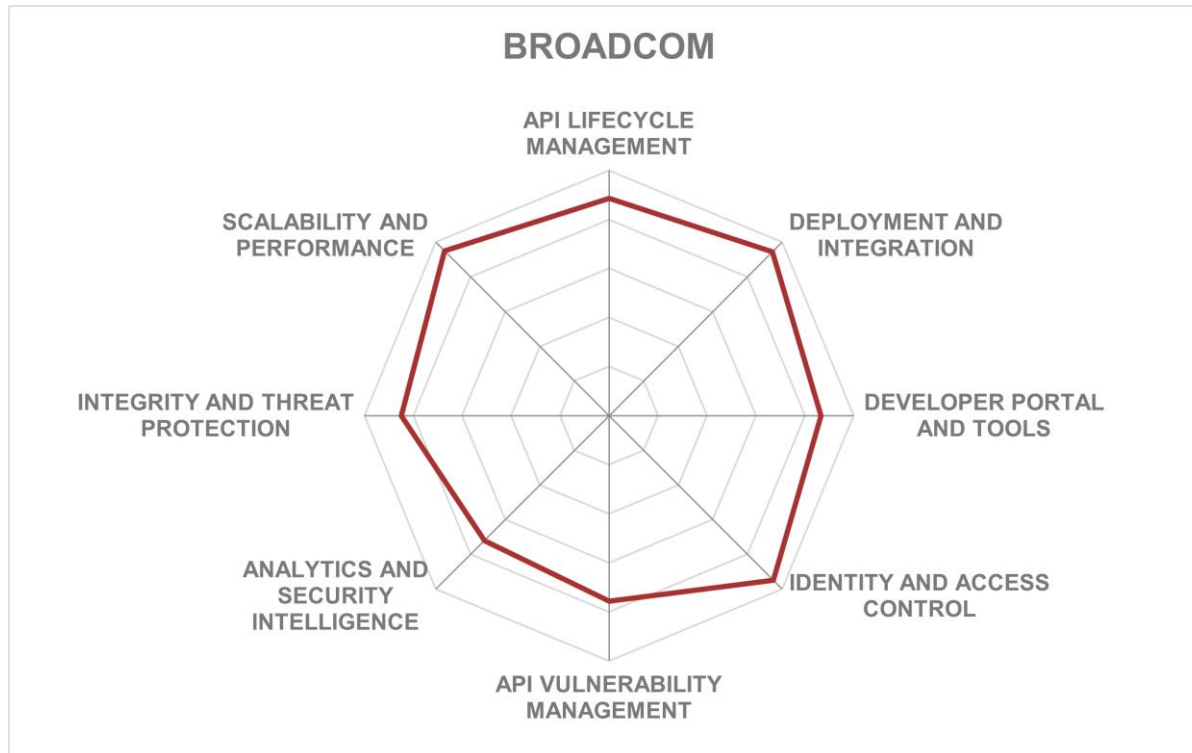
Strengths

- Federated API management across Axway and third-party gateways.
- Unified platform for API management, productization, and integration.
- Broad support for API protocols: REST, SOAP, GraphQL, gRPC, AsyncAPI.
- Advanced API discovery, including unmanaged APIs via 3rd party integrations.
- Visual API and integration design with no-code support in Amplify Fusion.
- Deep support for hybrid, multi-cloud, and containerized deployments.
- Rich analytics and observability, including prebuilt and custom dashboards.
- Comprehensive developer portal with multi-language and multi-marketplace support.
- Support for agentic AI protocols and workflows.

Challenges

- Targeted primarily towards large enterprise customers; might be too complex for smaller companies.
- Lacks native AI/ML-based threat detection; relies on partner integrations.
- Generative AI features are still under development and not broadly available.
- The built-in API firewall is limited, based on an open-source project.

Broadcom – Layer7 API Management



Broadcom Inc., originally established in 1961 and headquartered in San Jose, CA, is a global technology company known for its diversified portfolio of semiconductor and infrastructure software solutions. Following its acquisition of CA Technologies and subsequent expansion of its enterprise software business, Broadcom incorporated the Layer7 API Management brand into its portfolio. Layer7, with roots dating back to 2002 as one of the original API management innovators, now plays a key role in Broadcom's strategy for securing and managing digital infrastructures at enterprise scale.

The Layer7 API Management platform is positioned by Broadcom as a strategic solution for managing and securing APIs across hybrid, multi-cloud, and on-premises environments, with an emphasis on the broader role of APIs in enabling secure digital transformation. The platform is offered as an integrated suite covering the entire API lifecycle, with additional capabilities to support compliance, threat protection, and digital risk management. Broadcom's vision aligns API security with the broader IAM, DevOps, and operations portfolios, supported by educational outreach through the API Academy.

The suite includes several core components: the Layer7 API Gateway, Mobile API Gateway, Developer Portal, and Precision API Monitoring, all designed to work in concert. Together, they support the design, deployment, security, and operational governance of APIs in highly distributed infrastructures. The Developer Portal facilitates publishing and onboarding, the Gateway provides the runtime enforcement layer, and the management interfaces allow for both GUI-based and automated, CI/CD-integrated policy deployment. Extensive support for standards such as REST, GraphQL, gRPC, and SOAP, along with support for modern deployment models including Kubernetes and virtual appliances, enables the platform to adapt to varied enterprise requirements.

The product provides full lifecycle API management capabilities, starting with API contract design, mocking, and validation. Legacy modernization and protocol transformation are addressed with rich support for SOAP, XML-RPC, and custom formats, allowing users to expose legacy services as standard APIs. Policy templates offer configurable logic for authentication, threat protection, and routing, accessible through a low-code interface as well as a new policy-as-code option via the Graphman framework.

Discovery and inventory are enabled through a centralized Developer Portal that catalogs all APIs, their versions, policies, and compliance status. APIs can be tagged, classified, and assigned to usage tiers, with visibility and access governed by fine-grained role- and attribute-based controls. While API discovery of third-party or shadow APIs is not its primary focus, Layer7 can ingest API data from external registries and supports integration with service mesh frameworks to facilitate dynamic runtime discovery and routing.

Layer7 offers real-time protection against a wide range of OWASP API Top 10 threats and traditional web exploits. The gateway enforces strong authentication, supports OAuth2, OIDC, SAML, and certificate-based access, and offers a built-in STS for token translation. Policies can be customized to include content filtering, response shaping, and advanced routing logic. Rate limiting, quota enforcement, and access control policies are enforced per application, user, or organization, with multi-tenant support. The solution also includes native DDoS protection mechanisms, schema validation, and traffic anomaly detection through external AI/ML integrations.

Analytics and observability are addressed through the Layer7 Intelligence initiative, combining rule-based configuration checks with traffic analysis and AI-assisted insights. The platform supports OpenTelemetry and can export enriched traffic metadata to external SIEM, XDR, and analytics platforms. Auditing, alerting, and forensic capabilities are available out of the box, with support for both real-time and scheduled reporting.

Layer7 supports a flexible deployment model, including on-premises, private cloud, and container-based installations. The system is cloud-agnostic, supports any Kubernetes distribution, and includes a native operator for managing ephemeral gateways and GitOps-driven workflows. The API Gateway can be used as an AI Gateway, enabling secure access to LLM APIs, applying prompt filtering, SLA enforcement, and token usage metering. However, direct AI/ML capabilities within the product remain limited, relying on integrations with external platforms.

What sets Layer7 apart is its powerful, policy-driven enforcement engine, which supports flexible logic flows, protocol mediation, and extension through Java, JavaScript, and custom assertions. The platform's extensibility, combined with its proven ability to support high-scale, critical workloads across industries like finance, telecom, and public sector, makes it an attractive option for organizations with complex API strategies. However, the product's innovation pace in some areas has been relatively conservative compared to newer entrants in the market.

Layer7 API Management is best suited for large enterprises and government customers with stringent security, compliance, and deployment requirements. The platform, which serves many of the world's largest banks and telecoms, is optimized for regulated and mission-critical environments.

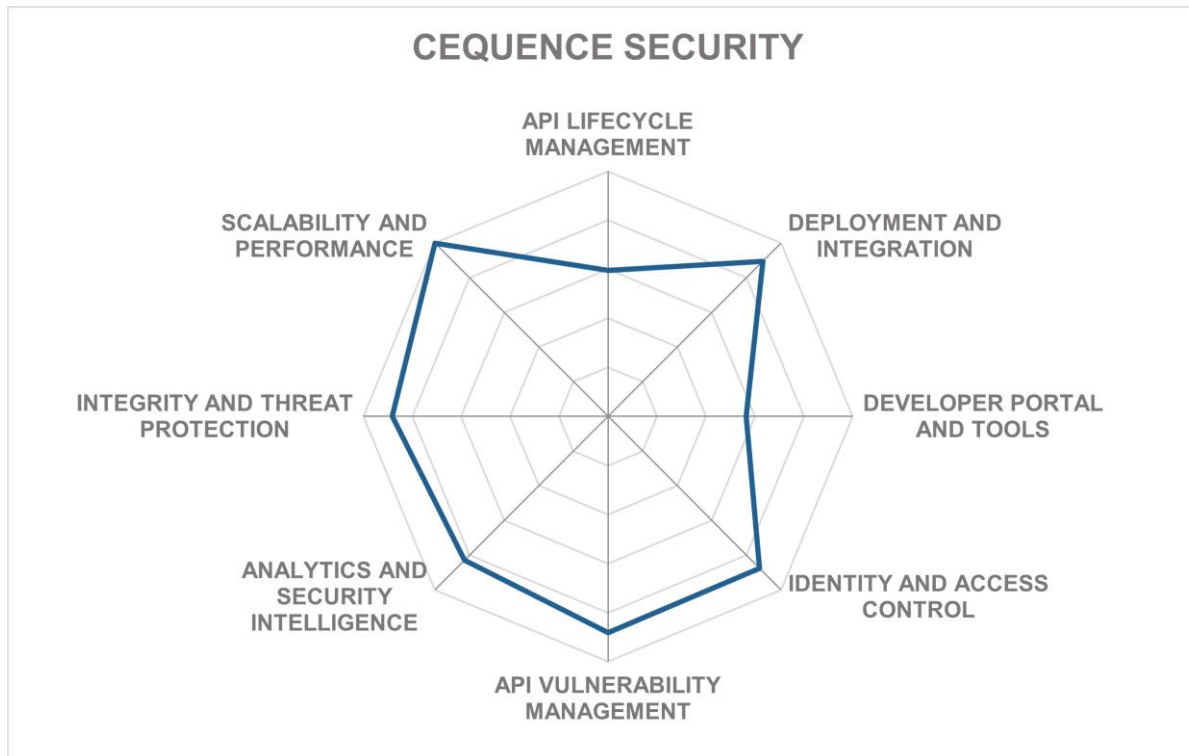
Strengths

- Full lifecycle API management from design to monitoring.
- Proven scalability across high-volume, mission-critical environments.
- Strong security policy engine with flexible enforcement options.
- Extensive protocol and format support, including legacy transformation.
- Cloud-native and hybrid deployment support with Kubernetes operator.
- AI Gateway capabilities for securing LLM applications.
- OpenTelemetry support for observability integrations.
- Advanced configuration and runtime intelligence features in development.
- Comprehensive developer onboarding and self-service via the portal.
- Highly customizable policies via low-code, policy-as-code, and SDKs.

Challenges

- Discovery of shadow APIs limited to integrations, not core functionality.
- No built-in dynamic API security testing or attack surface management.
- Limited native use of AI/ML for threat detection and traffic analysis.
- Relatively slow feature delivery in cutting-edge security innovation.
- Minimal out-of-the-box support for compliance frameworks.

Sequence Security – Unified API Protection



Leader in



Sequence Security, established in 2015 and headquartered in Santa Clara, CA, is a cybersecurity vendor with a focused mission to protect web applications and APIs from modern threats, including bot-driven abuse, fraud, and API-specific attacks. A team of seasoned security professionals built the company and developed a unique architecture centered on the Unified API Protection (UAP) platform. Today, Sequence protects well over 4 billion user accounts globally, positioning itself as a key player in the intersection of API security and bot management.

The company's strategic direction is built around avoiding any form of application or client-side instrumentation, a departure from many vendors who rely on JavaScript, SDKs, or in-app agents. Sequence instead embraces a network-based approach that allows for rapid, non-intrusive deployment. The goal is to reduce friction for application and security teams, enabling near-instant visibility, compliance checks, and real-time protection.

The Cequence UAP platform consists of three core components: API Security (focused on discovery, compliance, and testing), Bot Management (for real-time threat mitigation and fraud prevention), and API Spyder (for attack surface discovery). These are delivered as a tightly integrated suite with deployment flexibility spanning SaaS, hybrid, and fully on-premises models. Cequence's deployment strategy leverages a lightweight, high-performance data plane consisting of proxies (Defenders), sensors, and aggregators (Bridges), all deployable via containers or virtual machines.

API discovery is achieved through a combination of passive and active techniques. The system can ingest traffic via mirror ports or tap interfaces, analyze it using ML-powered heuristics, and classify APIs, whether internal, external, or third-party. API Spyder, the externally facing discovery module, complements this by analyzing DNS domains to identify exposed endpoints. Cequence also supports OpenAPI specification generation and validation, enabling proactive identification of schema drift, shadow endpoints, or undocumented APIs. Runtime traffic is continuously compared against specifications to identify non-conformance and policy violations.

The platform's API security posture management offers deep vulnerability assessment capabilities. Detected issues are aggregated and correlated across APIs and can be triaged with customizable workflows. These findings feed into automated remediation pipelines. Pre-production APIs can be tested with synthetic traffic via integrated tools, which can be triggered within CI/CD pipelines or manually, enabling shift-left security.

Analytics and threat intelligence are anchored in CQAI, the company's proprietary machine learning engine. It performs layered analytics including user fingerprinting, behavior profiling, and intent recognition. This allows the platform to identify both infrastructure-level and business logic attacks with high fidelity. Attack patterns are classified and labeled automatically using supervised learning, enabling actionable intelligence without manual intervention.

Threat protection is delivered in real time, both through inline proxies and out-of-band integrations. Cequence provides built-in defenses against transport-level threats, content-based exploits like SQL injections, and application-level DDoS. Its fingerprinting engine enables mitigation strategies that do not rely on fragile signals like IP addresses or user-agent strings. Instead, it evaluates payload characteristics, session context, and statistical anomalies. Defensive actions include blocking, deception, rate-limiting, or header injection.

Cequence is extending their capabilities to address AI-specific threats, including unauthorized usage of third-party LLMs by internal applications, as well as visibility into AI bots and crawlers. The platform can now identify traffic from known AI agents and detect unintended data exfiltration in prompts sent to external APIs like ChatGPT or Claude.

Differentiators of the Cequence UAP platform include its app-agnostic, zero-instrumentation approach and rapid deployment time. The company's native proxy-based enforcement capability and multi-layered behavioral analysis set it apart from first-generation API security vendors who rely heavily on WAFs or inline agents. Automated threat detection and policy creation powered by ML significantly reduce response times to emerging attacks.

The company's customer base spans North America, EMEA, Latin America, and Asia-Pacific, with notable deployments in regulated environments like the Middle East. Key use cases include runtime API protection, bot abuse mitigation, fraud prevention, API governance, and securing AI-driven applications. The product is particularly well-suited for organizations requiring sovereignty over data and infrastructure.

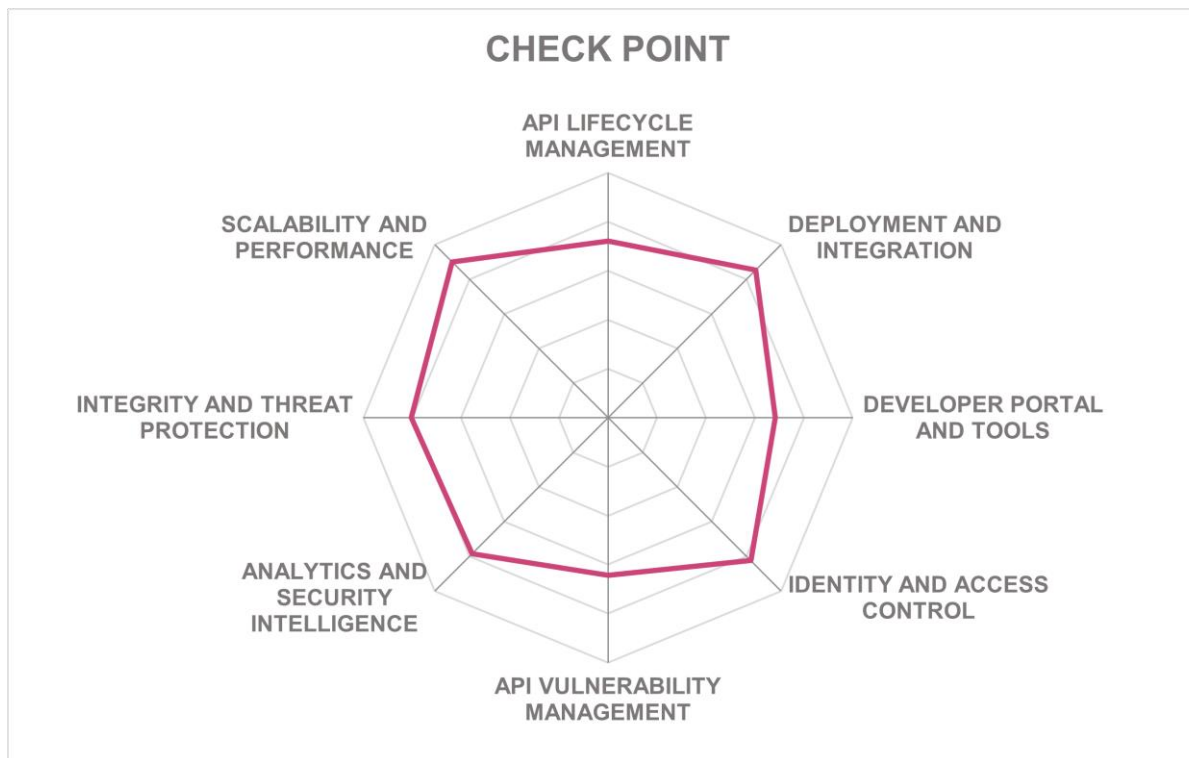
Strengths

- Unified platform integrating API Security, Bot Management, and Attack Surface Discovery.
- Support for full API lifecycle security, including pre-production testing.
- No client-side instrumentation or server-side agents required.
- Full support for passive and inline deployment models.
- ML-powered auto-discovery and inventory of internal, external, and third-party APIs.
- Real-time behavioral fingerprinting and intent analysis.
- Autonomous threat detection and policy enforcement.
- AI-specific threat coverage, including detection of LLM usage and AI bots.
- Flexible deployment across SaaS, on-premises, and hybrid environments.
- API virtual patching and shift-left testing support.

Challenges

- Pureplay API security solution, does not support API management capabilities.
- Focused primarily on HTTP APIs, limited support for event-driven protocols.
- Does not offer content transformation or protocol mediation features.
- AI security capabilities are in development, expected to be announced later.

Check Point – CloudGuard WAF



Leader in



Check Point Software Technologies, founded in 1993 and headquartered in Tel Aviv, Israel, is a long-established player in the cybersecurity industry. Its portfolio spans network, endpoint, cloud, and mobile security, unified under the Infinity architecture. Among its offerings, Check Point CloudGuard Web Application Firewall (WAF) serves as the company's solution for securing modern web applications and APIs. As a cloud-native tool built to scale across hybrid and multi-cloud environments, CloudGuard WAF aims to combine strong security with minimal operational overhead, reflecting Check Point's broader vision of prevention-first, AI-powered, and integrated security.

The company's approach to API security is based on three core principles: eliminating blind spots through deep and dynamic visibility, blocking threats proactively rather than reactively, and enabling centralized, automated management across diverse IT environments. CloudGuard is integrated into the broader Check Point Infinity platform, which enables it to leverage shared threat intelligence and policy orchestration capabilities, ensuring consistency across different layers of enterprise infrastructure. At the same time, the product

is designed to operate independently, either as a SaaS offering or as a flexible agent-based deployment using containers, Kubernetes, or virtual machines.

CloudGuard WAF offers native discovery and schema generation for all APIs in the environment, including undocumented and third-party endpoints. These are then used to enforce request validation and structural consistency, helping to block malformed or unauthorized traffic in real time. A standout feature of the product is its dual-layer AI engine. The first layer, referred to as the Attack Indicator Engine, is trained on millions of benign and malicious requests to identify multiple small indicators of attack rather than relying on matching signatures. The second layer adds contextual intelligence by analyzing the behavior of users, crowd patterns, trusted actors, and application-specific semantics. This multi-layered model contributes to the product's notably low false positive rate and strong performance in detecting even zero-day threats.

CloudGuard supports early API design and schema validation by observing live traffic and generating inferred schemas. These schemas can then be manually enriched or exported for external validation. Although the platform does not automatically infer complex business logic (such as value ranges or semantic constraints), it provides administrators with the means to define such rules through custom policies. This offers a flexible foundation for securing both internal and public-facing APIs, as well as legacy and modern application architectures. The platform also integrates with major CI/CD pipelines and supports declarative policy deployment, making it a good fit for organizations following DevSecOps methodologies.

The discovery and inventory capabilities are comprehensive, covering not only REST and GraphQL APIs, but also gRPC and other protocols when routed through supported gateways. CloudGuard can detect and classify APIs based on usage, traffic volume, and exposure status, and it maintains revision histories to track schema and behavior changes over time. For vulnerability management, the system correlates runtime behaviors with inferred schemas, applies threat intelligence from Check Point's ThreatCloud, and surfaces tuning suggestions based on model confidence levels. CloudGuard supports integrations with external testing tools like OWASP ZAP or Burp Suite.

Administrators can explore runtime dashboards showing high-risk APIs, schema violations, and anomaly trends. All traffic is subject to behavioral analysis and risk scoring. Detected incidents are mapped to MITRE ATT&CK tactics and can be exported to SIEM or SOAR platforms. The system supports JSON schema validation, protocol-specific protections, and rate limiting, including geo-aware traffic policies. Protection extends to common threats such as SQL injection and cross-site scripting, as well as OWASP API Top 10 risks. Notably, CloudGuard also includes features to safeguard AI workloads and APIs from LLM-specific attacks such as prompt injection.

The solution supports SaaS-based WAF as well as agent-based models that can operate as reverse proxies or Kubernetes ingress controllers. It can be embedded within or chained after third-party CDNs, integrated with service meshes like Istio or Envoy, and scaled dynamically using cloud-native mechanisms. Policy management supports grouping assets into zones, with "practices" (bundled sets of security behaviors) applied either centrally or

per asset. This zoning capability becomes especially important at scale, allowing customers to apply differentiated security levels across thousands of APIs.

CloudGuard WAF distinguishes itself through the maturity of its AI-driven detection engine and the comprehensiveness of its context-aware analysis. Unlike many rule-based or partially adaptive solutions, CloudGuard emphasizes preemptive prevention even for unknown attack variants. Its WAF comparison project claims to have a detection rate of 99.3% and a false positive rate under 1%, which is corroborated by public tests. The product's broad feature set makes it suitable for a wide range of customers, but it is particularly fitting for large enterprises operating in hybrid or multi-cloud environments with complex API ecosystems. At the same time, thanks to its SaaS model and automation-first design, it remains accessible to medium-sized organizations seeking enterprise-grade security without operational complexity.

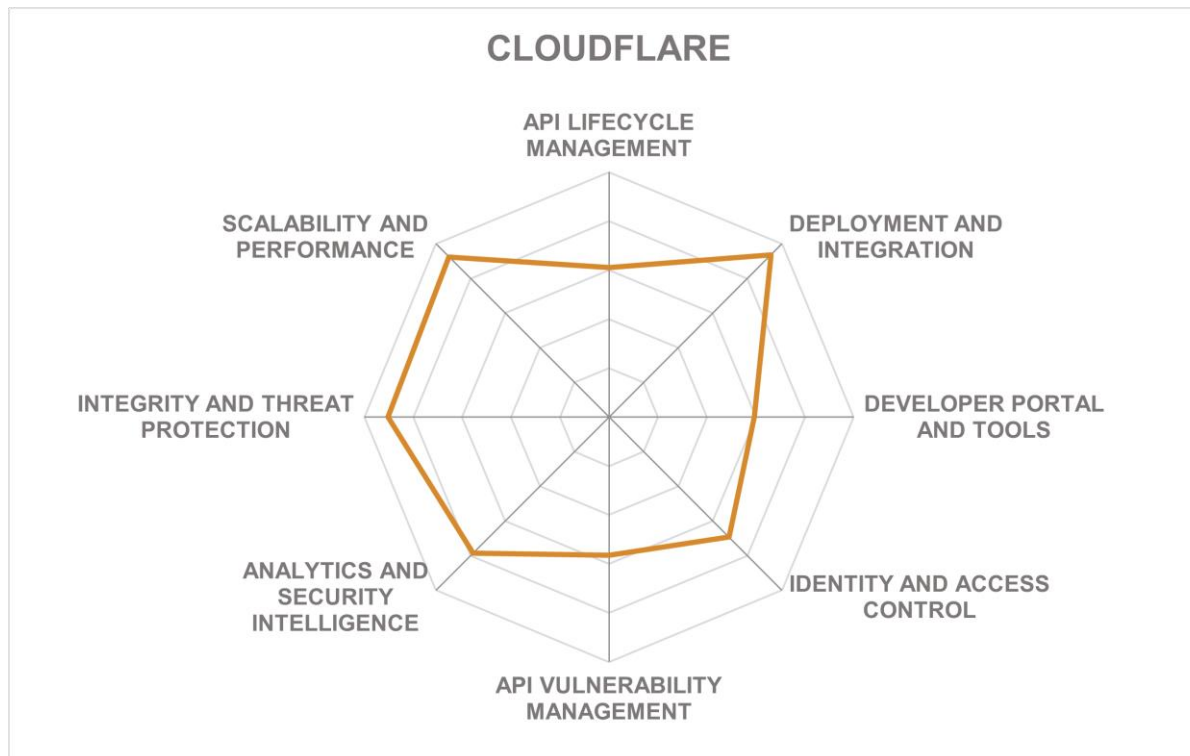
Strengths

- Strong API discovery, including third-party and shadow APIs.
- Broad support for different API standards and protocols.
- AI-driven threat detection using dual-layer machine learning models.
- Near-zero false positives combined with high detection rates.
- Contextual behavioral analysis using user, crowd, and application factors.
- Unified management through Infinity platform.
- Flexible deployment models including SaaS, agent-based, and Kubernetes.
- MITRE ATT&CK mapping and SIEM/XDR integrations.
- API-first architecture with full GraphQL-based management API.

Challenges

- Limited support for event-driven protocols like Kafka or MQTT.
- UI shortcomings can make management of large API inventories difficult.
- GenAI-assisted incident response still on roadmap.
- High reliance on cloud for advanced analytics may conflict with regulatory compliance.

Cloudflare – API Shield



Leader in



Founded in 2009 and headquartered in San Francisco, CA, Cloudflare is widely recognized for its expansive global network infrastructure and its comprehensive portfolio of application performance and security services. With presence in over 335 cities and interconnection with more than 13,000 networks, the company's platform is positioned at the heart of global internet traffic. Over the years, Cloudflare has steadily evolved from a content delivery and DDoS mitigation provider into a full-fledged security and performance platform. Within this broader context, API Shield, Cloudflare's dedicated API security solution, exemplifies its strategy of building security capabilities on top of a unified, globally distributed edge network.

The company has notably shifted focus from conventional API management to more security-centric capabilities in response to customer demand. The transition from the earlier "API Gateway" nomenclature to "API Shield" reflects this realignment. Leveraging the same infrastructure that powers its WAF, bot mitigation, and CDN offerings, Cloudflare seeks to offer a consolidated control plane for API security that is available at the network edge.

Thus, Cloudflare aims to differentiate itself from both traditional API gateways and specialized security vendors by delivering full-stack observability and protection without adding operational friction.

At its core, API Shield enables customers to discover, classify, and protect API traffic with minimal configuration. Capabilities include real-time traffic discovery, schema learning and validation, risk analytics, and volumetric abuse detection. Cloudflare provides inline request validation using uploaded or inferred OpenAPI schemas, integrated JWT verification, and rate limiting based on session-level context. Advanced features such as sequence analytics and mitigation, BOLA detection, and sensitive data labeling augment the platform's ability to detect complex threats and misconfigurations. Other services like WAF, DDoS protection, and bot mitigation further enrich the offering.

Once a customer onboards an API zone, the platform continuously analyzes traffic to identify endpoints, infer schemas, and classify usage patterns. Sequence analytics is used to model behavioral flows across API sessions, highlighting expected sequences and anomalies. These insights enable security features such as schema validation and sequence mitigation. Customers can augment discovered APIs with uploaded OpenAPI schemas. Endpoints can be labeled manually or automatically by Cloudflare, enabling filtering by use case, risk, or team ownership.

The platform's vulnerability management capabilities manifest in its posture management views and risk labeling system. Endpoints are continuously scanned for signs of missing or mixed authentication, sensitive data exposure, and suspected broken access controls. For example, Cloudflare's BOLA detection leverages behavioral entropy analysis to identify enumeration attacks at the session level. Recommendations are delivered through a centralized dashboard, with deep links into endpoint-level metrics and suggested remediations.

A unified console consolidates WAF, bot, and API security telemetry. Administrators can drill into request logs, analyze traffic volumes, investigate suspicious patterns, and apply mitigation rules. Risk scores, attack summaries, and actionable insights are prominently displayed in a posture overview page that combines alerts into a broader application landscape. Schema learning and validation also serve as foundational layers for building positive security models.

TLS and mTLS are fully supported, as is JWT validation with configurable claims-based access policies. Cloudflare's DDoS mitigation infrastructure is natively applied to API traffic, ensuring resiliency without any configuration. The platform's WAF capabilities, including ML-driven attack scoring and OWASP rules, are seamlessly available for API endpoints. GraphQL-specific protections, such as query depth and size enforcement, are included out of the box and benefit from Cloudflare's internal usage experience.

Cloudflare's architecture assumes a unique stance: rather than adapting to customer infrastructure, customers route traffic through Cloudflare's globally distributed edge. This simplifies operational complexity and ensures consistent policy enforcement across geographies, but it also means that protection is limited to APIs exposed over the public

internet. Nevertheless, the scalability, resilience, and manageability offered by Cloudflare's global footprint are compelling for many organizations.

While not a lifecycle management platform in the traditional sense, API Shield includes several developer-friendly features. Schema uploads, documentation generation, and hosted developer portals are integrated into the product. Support for real-time rule deployment, alert subscriptions, and deep API integration make the product more DevSecOps-oriented.

Among the product's most compelling differentiators are its ML-powered features that build on Cloudflare's unique traffic visibility. Sequence mitigation, a patented capability, allows organizations to enforce correct API usage patterns, a level of granularity rarely found in edge-based tools. Areas for improvement include broader protocol support, like gRPC or Kafka, and better integration with third-party gateways.

Cloudflare primarily targets enterprise customers with large, public-facing API estates. The product is also well-suited for organizations seeking to consolidate multiple layers of security into a single operational model, particularly those already using Cloudflare's WAF or CDN services.

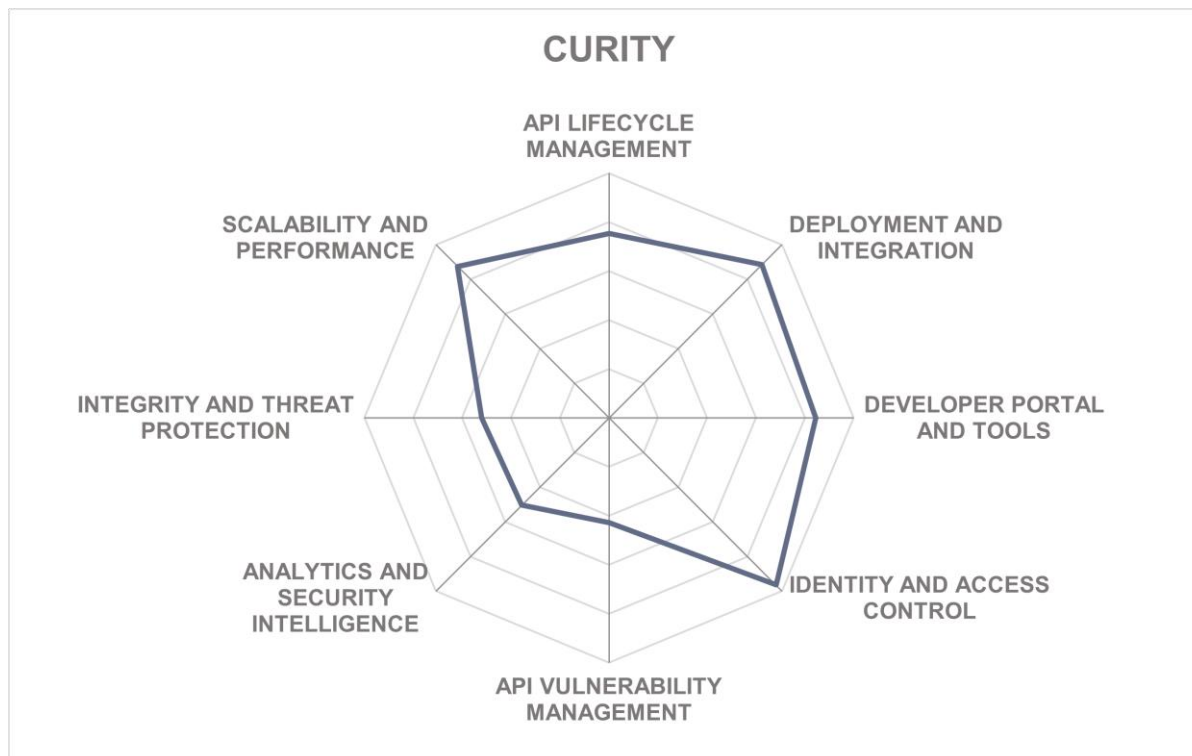
Strengths

- Seamless integration with Cloudflare's global edge infrastructure.
- Robust DDoS protection and bot mitigation at scale.
- Continuous API discovery with schema learning and validation.
- Unique sequence analytics and mitigation capabilities.
- Behavioral anomaly detection including BOLA scanning.
- GraphQL-specific protections included by default.
- Easy deployment and automation via CI/CD and Terraform.
- Unified console for web and API security analytics.

Challenges

- Limited capabilities for private/internal APIs not routed through Cloudflare.
- No native support for API design, testing, or full lifecycle management.
- Developer portal features are basic and lack full productization workflows.
- Protocol coverage outside REST and GraphQL is limited; gRPC and others are roadmap items.

Curity – Identity Server



Leader in



Curity, founded in 2015 and headquartered in Stockholm, Sweden, is an identity and access management vendor with a strong focus on API security. Built by a team of architects and identity specialists, Curity has established itself as a trusted partner for organizations seeking to modernize their digital infrastructure. The company's primary offering, Curity Identity Server, is designed as a cloud-native platform that combines IAM with advanced API protection capabilities.

The company's vision is to simplify and strengthen API security by providing a unified identity layer across distributed and hybrid application environments. By externalizing identity-related concerns and tightly coupling them with token-based access control, the Curity Identity Server enables developers and security teams to enforce consistent, policy-driven protection across heterogeneous systems and channels.

The product is built around a modular architecture with three core components: the Authentication Service, the Token Service, and the User Management Service. The solution

supports an extensive array of standards, including OAuth 2.0, OpenID Connect, SCIM, and Verifiable Credentials. Curity differentiates itself through strong support for customization, orchestration, and integration. With more than 30 built-in authenticators, a low-code orchestration engine, and the Hypermedia Authentication API, Curity enables dynamic and context-aware authentication experiences across web, mobile, and embedded clients.

Curity enables comprehensive API access control using dynamic tokens, token exchange, federated OAuth, and strong client authentication, such as workload identities. Tokens can be designed to include only context-relevant claims, reducing exposure while streamlining authorization logic in APIs. This approach supports both opaque and structured token formats and includes integrations with authorization engines such as OPA and emerging standards like AuthZEN. Curity's approach to securing application interfaces extends to mobile clients through a hardware-attested authentication flow, and to single-page applications as well.

Curity's support for DevOps and hybrid deployment scenarios is a notable strength. The solution can be deployed on-premises, in private or public cloud environments, and integrates natively with container orchestration platforms via Helm charts and infrastructure-as-code templates. Configuration-as-code, RESTCONF APIs, and CLI-based management ensure compatibility with modern CI/CD pipelines. Alarm management, audit logging, and telemetry integration with tools like Prometheus and Grafana provide observability into both operational and security-related events.

Where Curity stands out is in its deliberate choice to operate as a self-contained, deployable application rather than a SaaS offering. This gives customers full control over data locality, system residency, and compliance, which is increasingly important in regulated industries. The Hypermedia Authentication API, advanced claim management system, and native support for verifiable credentials position Curity well ahead of many competitors that offer only superficial identity federation or coarse-grained access control.

While its capabilities in securing and orchestrating identity flows are extensive, it does not provide API gateway functions or API lifecycle management features directly. The company relies on reference architectures and integrations with external API management platforms for those use cases. Similarly, the architecture can integrate additional security components that do additional jobs, such as AI-based threat detection, vulnerability management or policy management.

Curity's solution is most compelling for organizations with complex, high-assurance identity and API security requirements. It is especially relevant for digitally mature enterprises with a strong DevOps culture, multi-cloud architectures, and the need to implement custom or evolving authentication and authorization requirements at scale.

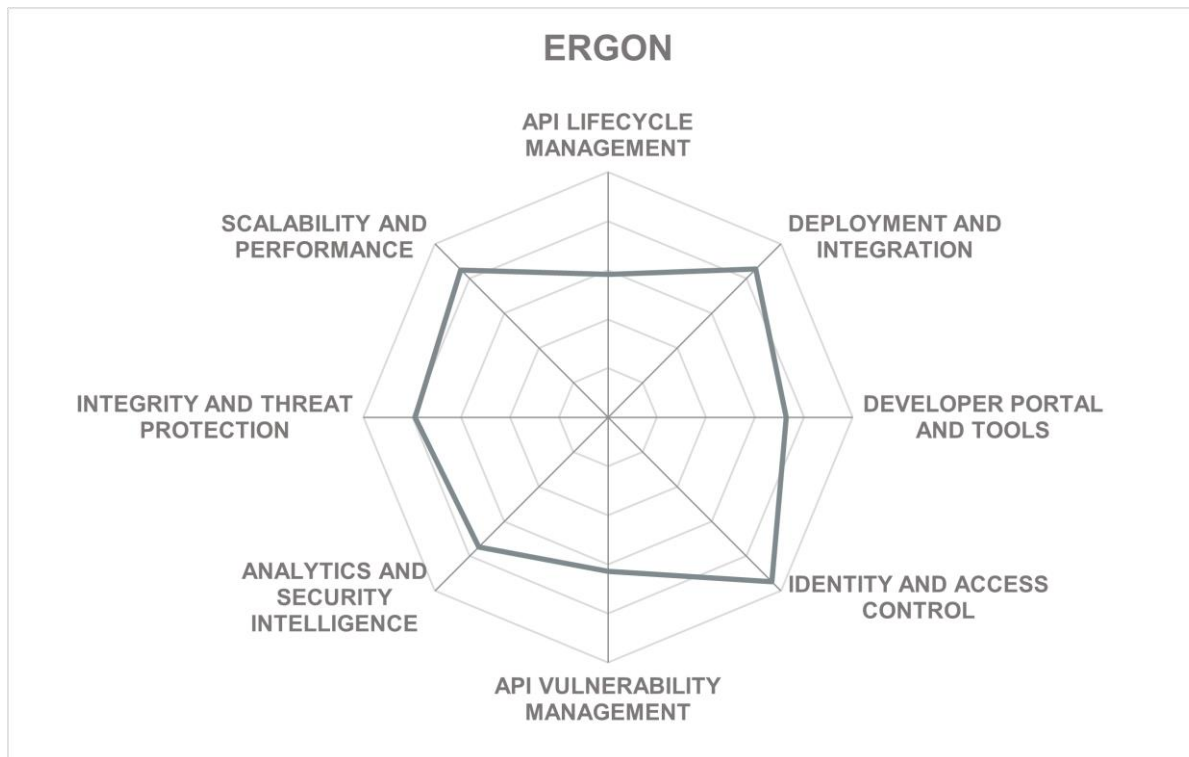
Strengths

- Deep integration of identity and API access control through standards like OAuth2, OpenID Connect, SCIM, and Verifiable Credentials.
- Highly flexible, low-code orchestration engine for dynamic user authentication journeys.
- Innovative Hypermedia Authentication API for secure, browserless native client interactions.
- Deployment in any environment: on-premises, private cloud, multi-cloud, Kubernetes.
- Built-in support for high-assurance authentication, including biometric, passkeys, and attestation.
- Strong DevOps support with RESTCONF API, CLI, and Helm charts for automation and scaling.
- Extensive observability and alarm management capabilities integrated with leading monitoring tools.

Challenges

- Does not provide a standalone API security solution and also requires the use of other security components.
- Relatively small market presence outside of Europe and the Middle East.
- Some administrative interfaces may require deeper technical expertise for effective use.
- Limited analytics and reporting capabilities without external SIEM or BI tooling.

Ergon – Airlock Secure Access Hub



Leader in



Founded in 1984 and headquartered in Zürich, Switzerland, Ergon Informatik AG is a long-established software engineering firm with a strong focus on cybersecurity. Its flagship product suite, Airlock Secure Access Hub, integrates identity and access management with comprehensive application and API protection. Serving customers across more than 30 countries, Ergon continues to expand beyond its DACH stronghold into regions such as the Middle East and APAC, where it is steadily building recognition.

Ergon's strategic vision centers on converging identity management and web application and API protection into a cohesive security fabric. Rather than treating access control and threat prevention as separate disciplines, Airlock unifies both domains. Ergon operationalizes this through continuous trust evaluation, real-time signal exchange between access and application layers, and decentralized enforcement architectures that scale from monoliths to microservices. By bridging identity, behavior analysis, and protocol-level protections, the company positions Airlock as an adaptive security mesh for digital services.

The Airlock Secure Access Hub combines several core components into a tightly integrated offering. Airlock Gateway delivers WAF and API security functions, available as a virtual appliance. Airlock Microgateway targets cloud-native environments, supporting sidecar and Kubernetes Gateway API modes. Airlock IAM is the identity backbone, available both on-premises and now as a SaaS offering. Collectively, these components provide advanced features such as token exchange, anomaly detection via ML models, fine-grained identity propagation, and federated identity support. The suite supports deployment across hybrid infrastructures, with flexible licensing options and integrations for DevOps and Kubernetes ecosystems.

Airlock's API protection starts with strong support for OpenAPI specification enforcement, including request and response validation, JSON schema whitelisting, and protocol-level filtering. Legacy backends can be fronted by Airlock through reverse proxying and identity augmentation. For continuous delivery pipelines, Airlock Microgateway is configured via Kubernetes CRDs and YAML, enabling seamless integration into DevSecOps workflows and test environments. Discovery and inventory of APIs remain outside of Airlock's current scope, and the solution does not integrate with third-party tools in this area. Consequently, API vulnerability management is likewise limited to what can be validated against declared API contracts and enforced policies at runtime.

For analytics, Airlock provides extensive logging, visual reporting via Kibana and Grafana, and out-of-the-box dashboards for usage, violations, and behavioral anomalies. Anomaly Shield, the machine learning module, supports automatic model retraining, transfer learning for immediate protection on new services, and session-based behavioral profiling. Attack types are categorized and linked to blocking actions such as rate limiting, session termination, and role downgrading, although mapping to frameworks like MITRE ATT&CK is not provided.

The solution offers robust integrity and threat protection, covering the OWASP API Top 10 comprehensively. It provides built-in defense against DoS, bot abuse, XML/JSON/SOAP attacks, and protocol misuse. Filtering can occur both centrally and at the edge, depending on the deployment mode. Integration with BrightCloud threat intelligence and optional ICAP-based filters for malware scanning add further layers of resilience.

Deployment flexibility is a major strength of the platform. Airlock supports traditional on-premises models, public and private cloud setups, and cloud-native use cases through containers and microgateways. SaaS versions of both IAM and Gateway components are in active rollout, and managed services are available via partners. The solution also supports hybrid trust architectures using token exchange to enforce access control across disparate security zones, without requiring changes to legacy applications.

Airlock's core differentiator lies in the synergy of access and application protection. Through deep integration of IAM and WAAP functions, the platform can propagate identity context through the stack, enforce authorization policies at runtime, and adapt defenses based on risk signals. Its ability to exchange identity tokens across trust domains and present them in backend-native formats is particularly notable. This minimizes friction for legacy systems while preserving Zero Trust principles. Other recent innovations include client behavior

analysis for bot detection, form-based anomaly models, and post-quantum cryptography support. These enhancements reflect Ergon's ongoing investment in foundational and future-proof security controls.

Airlock is well-suited for mid- to large-scale enterprises with complex hybrid infrastructures, particularly those in regulated sectors such as finance and government. With strong customer bases in the DACH region and growing presence in APAC and the Middle East, it is especially relevant to organizations seeking a unified platform for both access control and application/API protection with strict sovereignty requirements.

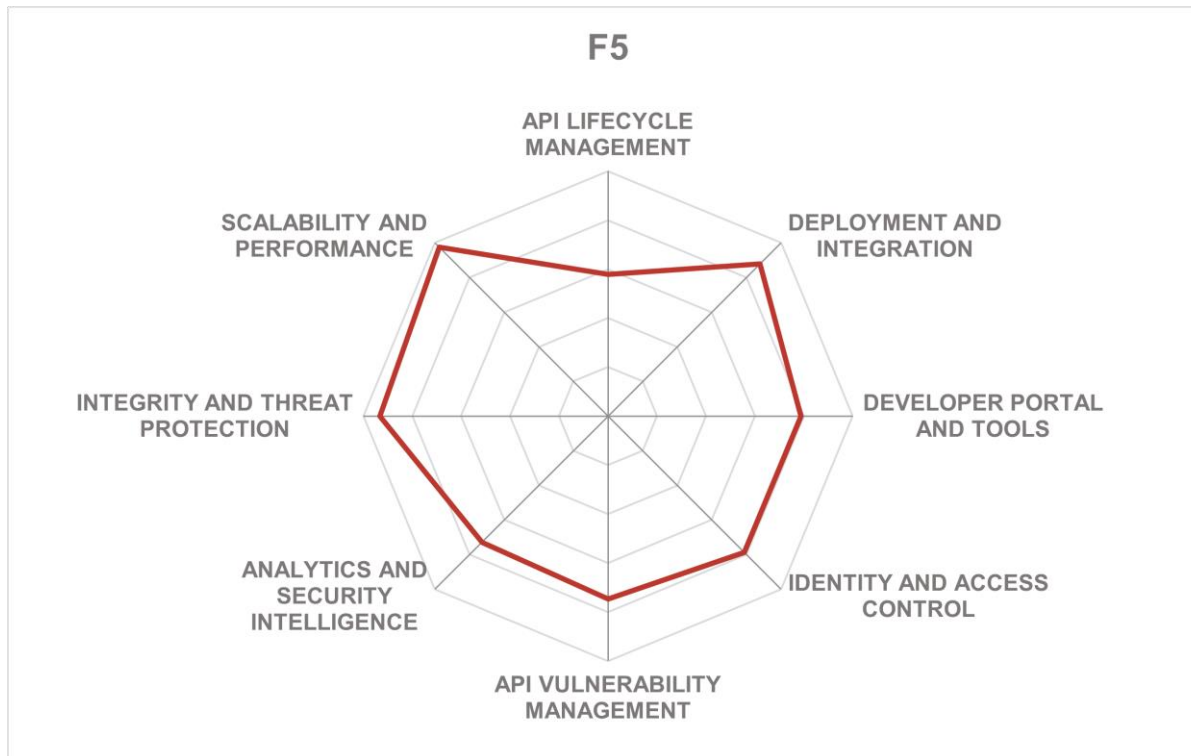
Strengths

- Fully integrated WAAP and IAM solution with unified signal exchange.
- Flexible deployment across on-premises, cloud, containers, and SaaS.
- Advanced identity propagation and token exchange support for Zero Trust.
- Strong machine learning-based threat detection via Anomaly Shield.
- Support for all OWASP API Top 10 threats and major protocol-level attacks.
- Built-in DoS, bot mitigation, and behavioral anomaly controls.
- Post-quantum cryptography PoC already validated in real environments.
- Comprehensive standard support (OAuth2, OIDC, SAML, FIDO2, etc.).
- Multi-language support and Airlock Academy for customer enablement.

Challenges

- Policy management for distributed deployments can be complex.
- Small but growing partner ecosystem and global market presence.
- No dynamic or proactive vulnerability testing or analysis.
- AI/ML capabilities do not currently include generative models.

F5 – Distributed Cloud API Security



Leader in



F5, established in 1996 and headquartered in Seattle, WA, is a veteran vendor in the application delivery and security space. Originally focused on load balancing and traffic management, the company has evolved its portfolio significantly lately, expanding into Web Application and API Protection (WAAP), bot mitigation, and security for modern cloud-native applications. With its Distributed Cloud Services platform, F5 aims to unify application delivery and security across heterogeneous infrastructures.

The company positions its Distributed Cloud API Security offering as a full-lifecycle API protection solution that brings together API discovery, runtime protection, threat detection, and compliance into a single control plane. This approach is underpinned by F5's extensive experience in application delivery and its architectural flexibility, which enables customers to

deploy security controls at any point in their infrastructure, from centralized SaaS consoles to distributed edge locations.

The solution combines several capabilities that are critical for modern API security. It supports multiple discovery mechanisms (traffic analysis, code repository scanning, and client-side crawling) to build and maintain a comprehensive API inventory. The platform then correlates this information to enable schema validation, anomaly detection, and risk scoring at the endpoint level. Protection rules can be defined and enforced inline, including rate limiting, IP reputation filtering, DDoS mitigation, and custom access controls. The solution also includes advanced features such as PII masking and AI/ML-powered behavioral analytics, all accessible through a unified SaaS-based dashboard.

It supports proactive testing of APIs during development through integration with source code repositories. This allows early identification of vulnerabilities before APIs reach production. Schema validation and dynamic discovery ensure that shadow and zombie APIs are identified and can be evaluated for compliance. Runtime security mechanisms include Layer 7 DDoS detection, automated mitigation, and correlation of signals from WAF, bot defense, and rate-limiting engines to flag malicious clients. The system supports real-time monitoring and analytics with configurable dashboards and graph-based views of API flows and endpoint-level metadata.

F5's platform can be deployed as a SaaS offering, through customer edge nodes, or in hybrid configurations, supporting a broad range of environments from Kubernetes and serverless to on-premises data centers. The architecture supports out-of-band API discovery from existing BIG-IP deployments, with additional integrations for NGINX and third-party gateways on the roadmap. The AI assistant within the console further simplifies operations by generating insights and remediation recommendations using natural language queries and retrieval-augmented generation.

What distinguishes F5 from many competitors is its platform-centric approach to API security, treating APIs as part of the broader application delivery fabric. Its support for GraphQL, SOAP, REST, and other protocols, and the ability to create protections without a single gateway or proxy provide significant architectural flexibility. The service also integrates well with DevSecOps workflows and supports CI/CD pipelines through its API-first design and Terraform provider. New capabilities are being implemented regularly, a notable recent release adding support for dynamic API security testing.

The product is particularly well-suited for organizations that already use F5 infrastructure and want to expand into modern API security without deploying a parallel security stack. Its hybrid architecture makes it attractive for customers with diverse and distributed IT estates, including those adopting AI workloads or seeking to secure LLM interfaces.

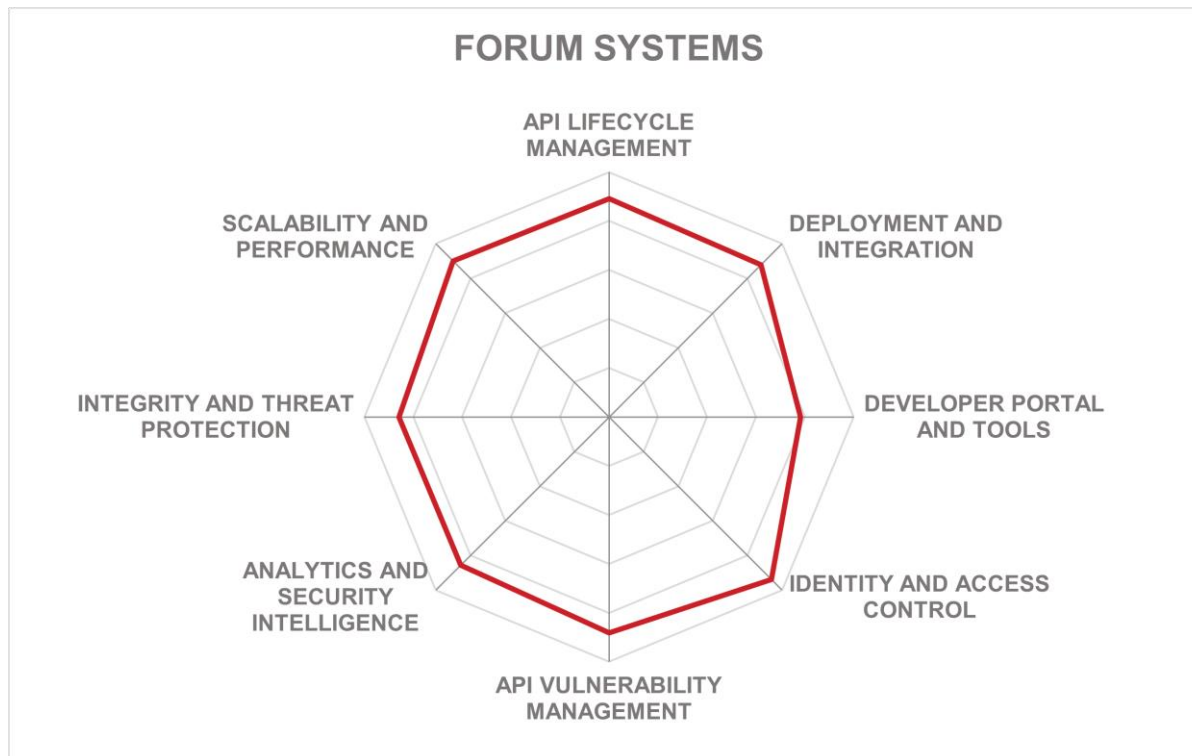
Strengths

- Broad API discovery coverage including code-based, traffic-based, and crawler-based methods.
- Comprehensive inventory and visualization features with endpoint-level insights.
- Integrated protection including schema validation, JWT inspection, rate limiting, and PII masking.
- Support for REST, GraphQL, gRPC, SOAP, and XML-RPC protocols.
- Out-of-band discovery integration with existing BIG-IP deployments.
- Unified SaaS-based management console with rich analytics and security dashboards.
- Flexible deployment models including SaaS, customer edge, and hybrid setups.
- AI/ML-powered behavioral analysis and contextual risk scoring.
- Built-in AI assistant for operational guidance and natural language queries.

Challenges

- Partial OWASP API Top 10 coverage; some detections remain roadmap items.
- Dependency on F5's own infrastructure for out-of-band discovery integration.
- GenAI protections are relatively new and not yet widely adopted.
- No support for event-based protocols such as Kafka or MQTT.

Forum Systems – Forum Sentry API Security Gateway



Leader in



Forum Systems, Inc., founded in 2001 and headquartered in Needham, MA, is a privately held company with a long-standing focus on secure API communication. Known for a steadfast commitment to “security by design”, the company has steadily evolved its flagship product, the Forum Sentry API Security Gateway, into a comprehensive and versatile platform. Forum Systems remains one of the few vendors in the space with a fully self-developed solution stack, avoiding third-party dependencies and external runtime libraries to maximize performance, resilience, and security integrity.

The company's strategic direction revolves around the convergence of traditional API security and the emerging needs of generative AI governance. Forum Systems emphasizes the unified role APIs play in both domains. With the addition of SecureGPT and the QS DevX developer experience platform, Forum has extended its value proposition to include LLM governance and API artifact lifecycle management into API and AI workflows. This

approach anticipates and addresses the architectural needs of agentic AI systems, where APIs serve as the critical interface for automation and data access.

Forum Sentry is offered as a standalone, fully integrated gateway platform, available in a wide range of form factors, including physical and virtual appliances, containers, cloud-native deployments, and as a managed service hosted in a HITRUST-certified AWS VPC. The product supports seamless upgrades with hot-swap capabilities, autonomous deployment pipelines, and unified policy management across deployment scenarios. These features are particularly appealing to enterprise customers in regulated sectors such as healthcare, government, and finance.

At its core, Forum Sentry delivers a broad set of capabilities across the entire API security lifecycle. The platform supports policy-driven API design and transformation, including legacy protocol mediation, schema validation, and format conversions. It provides native OpenAPI, AsyncAPI, and gRPC support. QS DevX enables lifecycle artifact management, credential control, GitHub integration, documentation versioning, and even AI model fine-tuning dashboards for SecureGPT environments.

API discovery and inventory management are facilitated through customizable templates and Backstage catalog integration, allowing centralized registration of all APIs deployed across gateway instances. The platform also offers dependency mapping and catalog graph visualizations, helping users understand the interdependencies among services. However, while internal and native APIs are well-covered, third-party API discovery is currently unsupported, which may limit visibility in more diverse hybrid environments.

Forum Sentry provides built-in static and dynamic security analysis for identifying API vulnerabilities and misconfigurations. These include schema and signature validation, rate limiting, payload inspection, and conformance to the OWASP API Security Top 10. The gateway also includes features for managing API attack surfaces, including XML and JSON firewalling, support for both symmetric and asymmetric cryptographic operations, and content extraction for threat detection. Advanced caching mechanisms and detailed transaction latency metrics enhance performance and protection against abuse.

The platform delivers security analytics and intelligence via machine-learning-ready metadata logging format. Forum's AI logs capture granular metrics per transaction, enabling streamlined ingestion into analytics platforms. Real-time dashboards and visualizations of API and model usage are available in QS DevX, and integrations with SIEM solutions are supported. The logs are cryptographically signed to ensure tamper resistance and audit integrity, supporting forensic investigations and compliance reporting.

Forum Sentry supports virtually all modern environments, including Kubernetes, hybrid cloud, on-premises, and managed SaaS scenarios. The solution handles protocol and message translation across a broad array of standards and legacy messaging formats, making it particularly useful for enterprises undergoing modernization without full rip-and-replace initiatives. Autonomous policy provisioning and agentless identity enforcement reduce operational complexity and improve security posture across distributed architectures.

The company has recently launched SecureGPT, a multi-LLM governance gateway designed for AI risk mitigation. It includes its own moderation LLM for prompt filtering, obfuscation of sensitive content, observability, and role-based access guardrails. It is offered both as a managed cloud service and as an on-premises appliance, making it viable for regulated environments. Its convergence with Forum Sentry under a unified architecture and UI reflects a forward-looking strategy to support agentic AI systems.

Forum Systems continues to differentiate itself through its engineering-centric, secure-by-design approach, offering high-performance capabilities without reliance on external runtime components. Notable differentiators include the absence of per-user or per-API licensing in instance deployments, advanced protocol mediation capabilities, and a uniquely integrated AI governance solution. However, areas for improvement include support for more AI-assisted features within API workflows, and broader support for external API discovery.

The company serves primarily large enterprise and government customers, with a customer base concentrated in North America and expanding presence in EMEA and APAC. Forum Sentry is particularly appealing for organizations seeking secure protocol mediation, AI governance integration, and performance optimization within hybrid or containerized infrastructures.

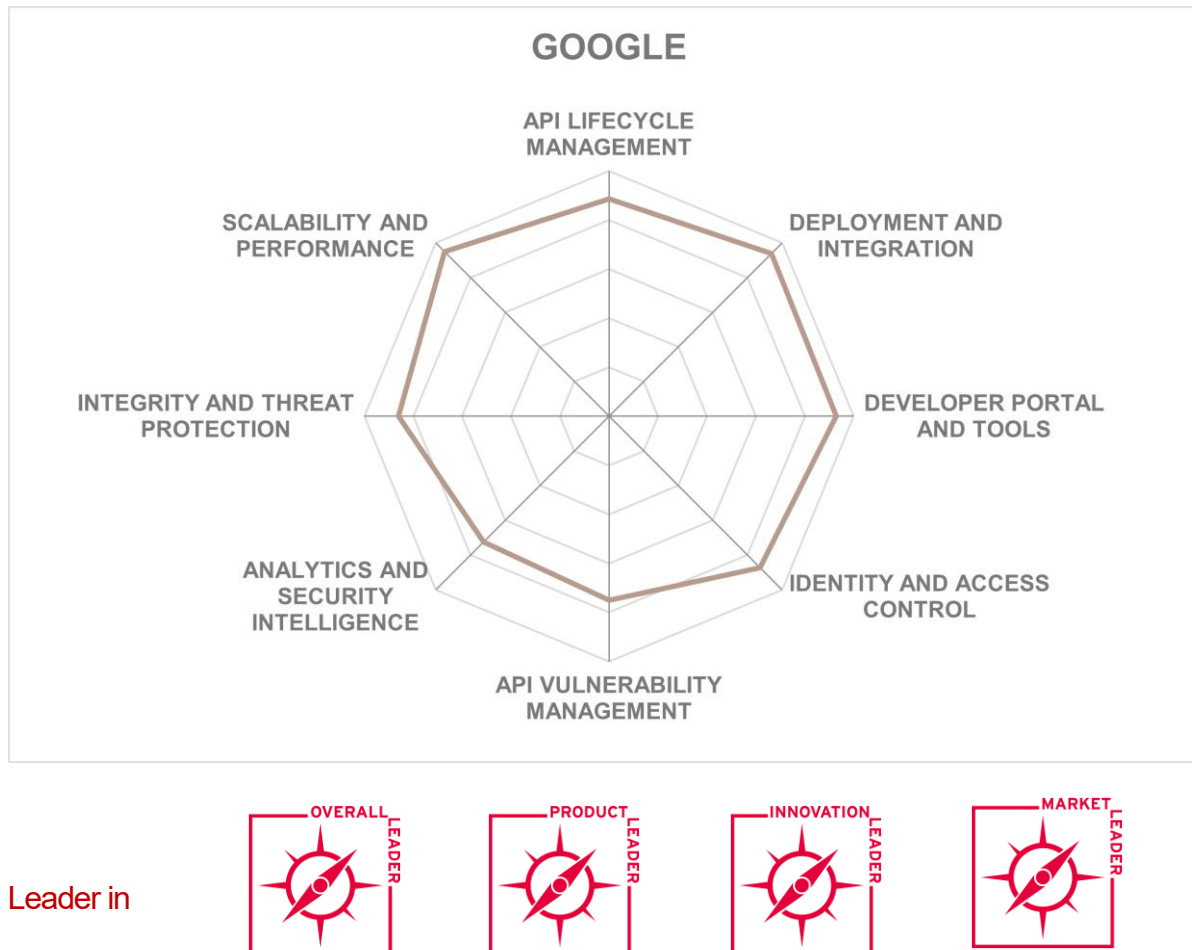
Strengths

- Security-by-design architecture with no third-party dependencies.
- Full range of deployment options including physical, virtual, containerized, and SaaS-managed.
- Autonomous Git-based deployment and silent installation workflows.
- Comprehensive protocol mediation and legacy modernization capabilities.
- Extensive cryptographic support including dynamic PKI, encryption, redaction, and signature validation.
- Unified developer experience via QS DevX platform with GitHub and Backstage integration.
- Integrated GenAI governance platform with obfuscation and prompt filtering.
- Full support for FIPS 140-2 and NIAP certification for high-assurance environments.

Challenges

- Limited support for discovering or managing third-party APIs.
- Lack of AI/ML-based threat detection in core API security functions (roadmap item).
- AI governance capabilities are only available through a separate (yet integrated) solution.
- Integration with service meshes is not currently supported.

Google – Apigee API Management



Apigee, now part of Google Cloud since its acquisition in 2016, is headquartered in Mountain View, CA. Founded in 2004 and a key contributor to the OpenAPI Initiative, Apigee has evolved into one of the most mature and comprehensive platforms in the API management space. Now fully embedded within Google's cloud services ecosystem, Apigee continues to expand its capabilities and integration depth, aiming to serve both enterprise and cloud-native customers with a broad, scalable, and intelligent API management solution.

Apigee's strategic vision is centered on delivering end-to-end visibility, governance, and control over APIs in increasingly complex IT environments. A key differentiator for Google's strategy is the platform's tight integration with other Google Cloud offerings such as Cloud Armor for threat protection, Cloud IAM for fine-grained access control, and Vertex AI for generative AI capabilities.

The Apigee API Management platform provides robust capabilities for designing, securing, publishing, monitoring, and monetizing APIs. It supports multiple deployment models including public cloud, hybrid, and private cloud, thus addressing the needs of enterprises

operating in regulated or hybrid environments. With Apigee X, the most recent generation of the platform, users benefit from deep integration with Google Cloud's AI-driven services, enhanced automation for API operations, and improved analytics. Apigee also includes an API hub that acts as a centralized repository for API discovery, documentation, and governance. Integration with Gemini Code Assist helps developers interact with the platform using natural language to generate API specifications and configure policies.

API discovery and inventory management are addressed through the API hub and associated cataloging tools, which allow organizations to maintain a unified and continuously updated registry of APIs. The API hub supports tagging, metadata enrichment, and lifecycle tracking, enabling better classification and governance of services across hybrid environments and for use with AI agents. Vulnerability management is delivered through automated conformance checks and posture assessments, powered in part by AI. These tools analyze API specifications and traffic to identify discrepancies, misconfigurations, and deviations from best practices. Security issues, such as improper authentication or overly permissive access controls, can be flagged early in the development process.

Apigee's analytics and monitoring functions are integrated with Google Cloud's broader observability stack. These tools provide real-time visibility into API traffic, usage trends, and anomalies. Security intelligence features can detect abnormal patterns, suspicious behaviors, and emerging threats. Integration with SIEM and XDR platforms allows these insights to be correlated with other infrastructure and application security signals. Apigee complements Cloud Armor (or the customer's WAF of choice) by providing fine-grained API runtime protection. It implements dynamic checks and alerting for API security misconfigurations, and ML- and heuristics-based rules to detect API attacks, even among traffic with valid credentials. It supports OWASP-compliant security policies to guard against common API threats such as injection attacks and credential stuffing and implements mutual TLS, rate limiting, quota enforcement, and IP filtering, among other controls. API keys, OAuth 2.0, and JWT-based authentication are supported out-of-the-box, enabling fine-grained access control tailored to enterprise requirements.

Support for hybrid and multi-cloud environments remains a core feature of Apigee's value proposition. The platform can be deployed in Kubernetes-based environments, on-premises, or across multiple public clouds. Through integrations with Istio and Envoy and an API management operator for Kubernetes, Apigee provides native support for microservices and service mesh architectures, enabling consistent policy enforcement and observability regardless of deployment topology.

Compliance support is broad, with tooling and templates available for meeting regulatory requirements such as GDPR, HIPAA, and others. Policy enforcement points can be customized to implement data residency, access control, and auditing mandates specific to industry verticals. Reporting and audit logging capabilities, when used in combination with Google Cloud's compliance offerings, enable easier adherence to both technical and legal standards.

Apigee integrates natively with Google Cloud services and provides connectors for third-party security solutions, developer tools, and IAM platforms. Apigee API hub also includes a

toolset for Agent Development Kit, so that developers building custom agents can give those agents access to their governable and secure enterprise APIs. Google is actively integrating protections for AI-based services and APIs, including guidance on preventing injection attacks and data leakage in LLM applications. Apigee's integration with Model Armor and its native AI policies help customers extend their API functionality to generative AI workloads deployed on Vertex AI or third-party platforms. This further demonstrates the vendor's forward-looking posture on securing emerging digital interfaces.

The core differentiators of Apigee include its native integration with Google Cloud's extensive service portfolio, which enhances its scalability, automation, and security capabilities. However, the platform's complexity and breadth may present a steeper learning curve for smaller teams or organizations lacking Google Cloud expertise. Apigee is best suited for large enterprises and regulated industries that require strong governance and comprehensive control over their API ecosystem. It is particularly valuable for organizations already invested in Google Cloud, as well as those building AI-enhanced or hybrid cloud applications.

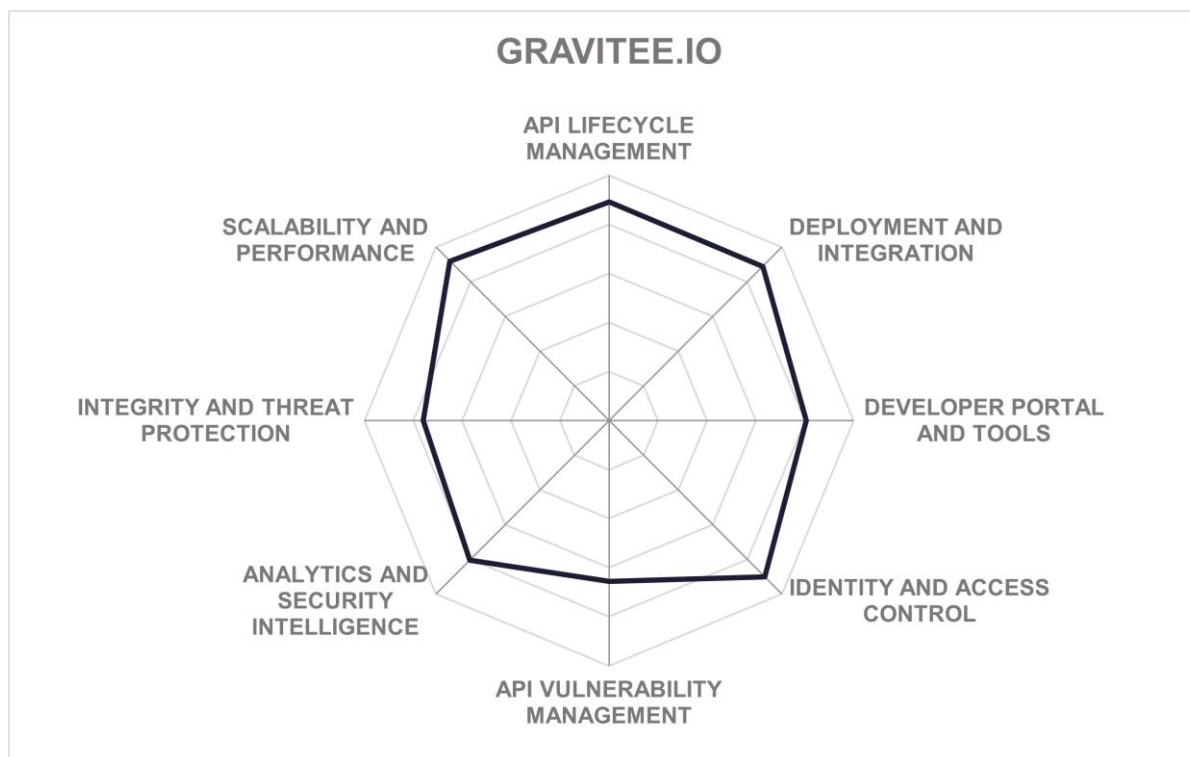
Strengths

- Native support for REST, SOAP, GraphQL, and gRPC.
- Deep integration with Google Cloud services including IAM, Model Armor, Cloud Monitoring, and Vertex AI.
- API hub for discovery, documentation, and governance.
- Duet AI assistant for natural language-based development and configuration.
- Advanced posture assessment and conformance checking.
- Robust analytics via Apigee Sense and Google Cloud monitoring tools.
- Extensive hybrid and multi-cloud deployment support.
- Native service mesh integration with Istio and Envoy, and a native Kubernetes operator.
- Built-in API and AI monetization capabilities and developer portal support.
- AI security enhancements for LLM and generative API workloads.
- Strong compliance and auditing support across verticals.

Challenges

- Steeper learning curve for teams unfamiliar with Google Cloud.
- Integration with non-Google services is possible through APIs, with additional effort.
- Complex configuration options may overwhelm small or mid-sized teams.
- Advanced AI features are still evolving and may require tuning.

Gravitee – API Management Platform



Leader in



Founded in 2016 and based in Boulder, CO, Gravitee has emerged as a prominent force in the API management space. Initially established by a group of developers in France, the company has since expanded its operations globally, while continuing to develop and maintain its open-source roots. Gravitee remains committed to redefining API management with a clear vision: unifying the governance, lifecycle, and security of both synchronous and asynchronous APIs in a single platform.

Gravitee's strategy focuses on providing a developer-friendly platform that supports REST, GraphQL, SOAP, WebSocket, Kafka, and MQTT equally. This integrated approach enables organizations to manage and secure various types of APIs and event streams through a single solution. Gravitee positions itself as an “event-native” API management provider,

extending lifecycle management, policy enforcement, and access control capabilities across all major integration styles. More recently, the company has started to position its platform in the broader “agentic API and event management” space, focusing on support for emerging use cases involving AI agents and agent orchestration protocols.

The Gravitee API Management Platform provides comprehensive lifecycle capabilities through a modular architecture. Its key components include Gravitee Cockpit for centralized administration across environments, Gravitee API Management for runtime governance, Gravitee Access Management for authentication and authorization, and Gravitee API Designer for visual, low-code API modeling. These are complemented by tools such as the Gravitee Alert Engine and a customizable developer portal. The platform supports a wide range of deployment models from fully on-premises to hybrid and fully managed SaaS, with native Kubernetes support.

API publishers can define and model API specifications using a visual interface in API Designer, which supports OpenAPI standards and includes versioning and promotion workflows. Policies for throttling, quota management, schema validation, and access control can be defined using a visual Policy Studio. This enables teams to define traffic plans tied to authentication mechanisms like OAuth2, JWT, or API keys, as well as more advanced flows such as mutual TLS or HMAC-based authentication.

Gravitee's federation capabilities aggregate APIs from other platforms, such as Amazon API Gateway, Apigee, or Confluent, into a unified developer experience. Discovered APIs are imported and cataloged automatically, with metadata enrichment and classification capabilities available in the management UI. Gravitee's governance model includes a rule-based scoring system that enforces best practices, including API freshness, documentation completeness, and authentication plan coverage. These scores are dynamically calculated and can be integrated into CI/CD pipelines to enforce promotion policies.

The platform collects performance metrics, usage statistics, and security-related events, which are visualized through customizable dashboards. Users can create and export reports or integrate with external SIEM and observability tools such as Datadog and Elasticsearch. Anomaly detection and alerting support real-time notifications or webhooks based on configurable thresholds. Administrative activities leave full audit trails, and all logs are protected against tampering.

Gravitee supports all modern identity standards, including OAuth2, OpenID Connect, SAML, and Kerberos, and offers built-in user federation and step-up authentication mechanisms through Gravitee Access Management. For API traffic, policies can be applied at the request, response, or message-frame level (for WebSocket or Kafka payloads). Protocol-specific protections are available for JSON, XML, and Avro. A wide range of threats including OWASP API Top 10, SQL injection, and session hijacking are addressed natively through configurable policies. Though native DDoS mitigation is not included, spike arrest and caching can be combined with third-party defenses.

Gravitee supports a wide range of environments, including Kubernetes, hybrid cloud, private and public clouds, and even serverless components. Gateway deployment is flexible and

scalable, with tagging and sharding support for managing geographically distributed clusters. APIs can be deployed declaratively or through a rich RESTful management API and GitOps-style workflows.

AI-related capabilities have recently been introduced. Gravitee now provides an AI gateway module that supports traffic shaping based on LLM token usage, prompt filtering, and even prompt security to block toxic or jailbreaking content. The platform can transform APIs into MCP-compliant tools, enabling integration with AI agents across different ecosystems. Additionally, Gravitee is building an “agent mesh” to apply governance, scoring, and lifecycle controls to LLM agents, like it does for APIs.

Among Gravitee's notable differentiators are its native support for event brokers like Kafka and MQTT, its ability to expose event streams through web-friendly protocols (such as WebSocket), and its groundbreaking Kafka proxy gateway that enforces API-style access policies on native Kafka traffic. However, the platform currently lacks native dynamic security testing or API attack surface management.

With comprehensive support for both legacy and modern protocols, Gravitee is particularly relevant for enterprises undergoing digital transformation with heterogeneous infrastructures or those adopting event-driven architectures.

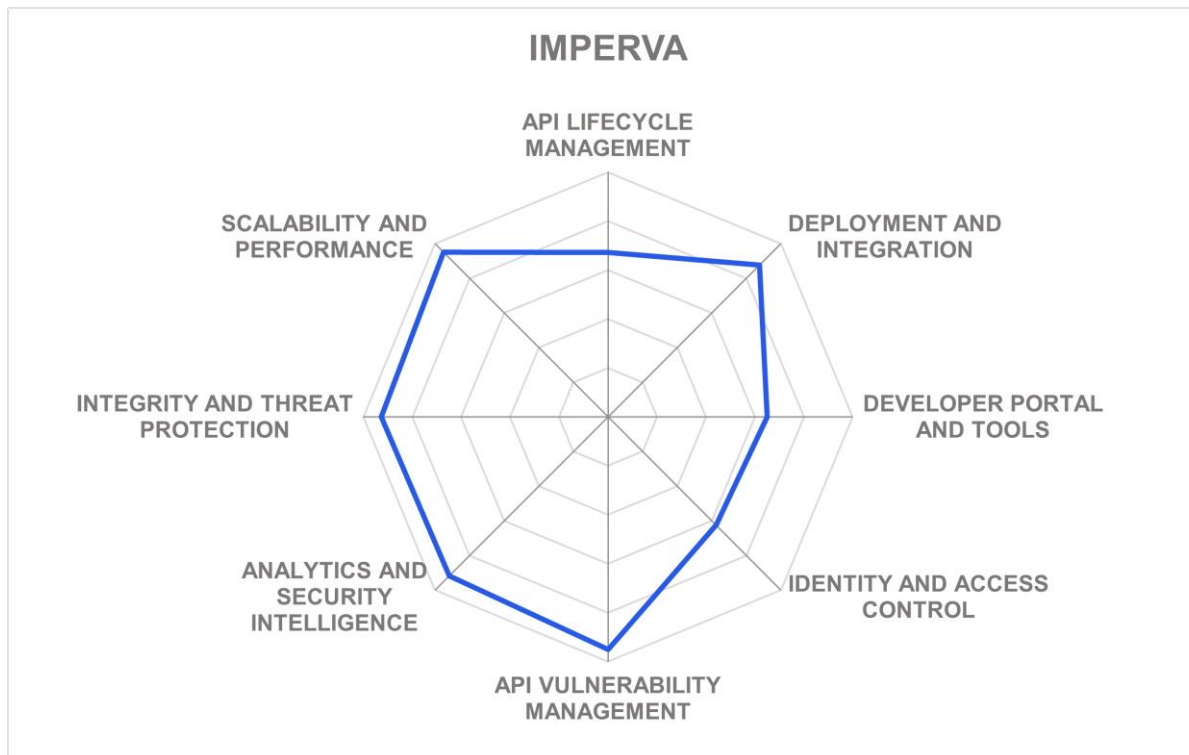
Strengths

- Unified platform for synchronous and asynchronous APIs.
- Native support for Kafka, MQTT, WebSocket, and other event protocols.
- Kafka gateway with policy enforcement on native Kafka traffic.
- Federation of third-party gateways and brokers into a unified portal.
- Visual API and policy design interfaces with support for Spectral rules.
- Strong integration capabilities (SIEM, DevOps, CI/CD, IDPs).
- Flexible deployment options with Kubernetes-native tooling.
- AI gateway support for LLM-specific use cases and prompt protection.

Challenges

- No built-in API attack surface management or dynamic security testing.
- Relatively low market visibility in North America compared to larger players.
- Generative AI and agent governance features are still a work in progress.
- Policy definition interface is technically oriented and may not suit non-IT users.

Imperva – Application Security



Leader in



Imperva is a cybersecurity vendor founded in 2002 and headquartered in Redwood Shores, CA. Known initially for pioneering the Web Application Firewall (WAF) space, the company has expanded its portfolio to cover a broad spectrum of data and application security challenges. With the acquisition of CloudVector, Imperva has evolved into a mature, full-stack application security vendor with a comprehensive API security solution at the core of its offering. After being acquired by Thales Group in 2024, Imperva continues to build on its vision of securing applications and APIs in any deployment model across hybrid and multi-cloud infrastructures.

Imperva's strategic direction in API security reflects a clear response to enterprise demands for integrated, consistent protection across fragmented IT landscapes. Their vision aims to deliver a unified application and API protection engine that can be deployed across any infrastructure, whether legacy, hybrid, cloud-native, or container-based. Imperva focuses on

interoperability and operational simplicity. Their solution is positioned to serve security teams that need to work alongside DevOps and API platform owners, without imposing restrictions on their development or delivery pipelines.

Imperva Application Security is delivered as a fully integrated SaaS platform with optional on-premises and hybrid deployment models. The solution is built around the Imperva Security Engine that connects detection and enforcement layers across diverse ingress points, including cloud WAFs, load balancers, Kubernetes ingress controllers, and third-party API gateways. For legacy or air-gapped environments, the same engine can be deployed on-prem via Imperva's WAF Gateway. All environments are controlled through a unified console, ensuring consistent management and incident response.

The platform enables full API lifecycle visibility starting with automated API discovery that supports shadow, zombie, and unauthenticated endpoints. Data classification is integrated directly into the discovery workflow, allowing security teams to identify and prioritize high-risk APIs handling sensitive information. Schema enforcement and OpenAPI validation support a positive security model that blocks requests deviating from defined specifications. These protections are augmented with behavior anomaly detection and contextual risk scoring to deal with threats like broken authentication, excessive data exposure, and logic abuse in real time.

Runtime protection extends to detection of business logic abuses such as Broken Object Level Authorization (BOLA) and Broken Function Level Authorization (BFLA). This is complemented by static posture analysis that enables customers to identify risky configurations. Security controls can be applied both via declarative policies across API classes and granularly, using parameter-level logic. Customers can block unauthorized calls to deprecated APIs, enforce schema protections, or activate bot prevention policies with a single click.

Imperva also integrates advanced analytics and threat intelligence capabilities through its Attack Analytics module. This system correlates and prioritizes thousands of individual events into incidents, delivering detailed insights into actors, tools, and techniques observed in API-related attacks, as well as automated policy recommendations. The platform maintains secure processing with anonymization and data residency controls for regulatory support.

The solution's flexibility shows both in support for all major API protocols including REST, GraphQL, gRPC, and SOAP, and deployment across Kubernetes, VM-based, and third-party API management environments. Imperva's AI and ML functions are applied for proactive detection and for automating policy generation and response actions. Furthermore, it is actively expanding those capabilities to protect GenAI-enabled applications, treating prompt injections and LLM abuse as natural extensions of its API security framework.

While many competitors depend on ecosystem integrations to perform enforcement, Imperva offers native blocking capabilities across environments. This, combined with its extensive patent portfolio and large-scale deployment flexibility, positions the platform well for organizations with complex, distributed IT infrastructures. However, while Imperva can

integrate with API gateways and some development tools, it does not provide SDKs, virtualization tools, or testing environments directly tailored for development teams.

Imperva's Application Security solution is best suited for large enterprises and regulated industries that operate hybrid or multi-cloud environments and require consistent, scalable security across both web applications and APIs. Its comprehensive threat prevention, inline enforcement, and unified management capabilities make it particularly attractive to security teams developing their consolidated operations.

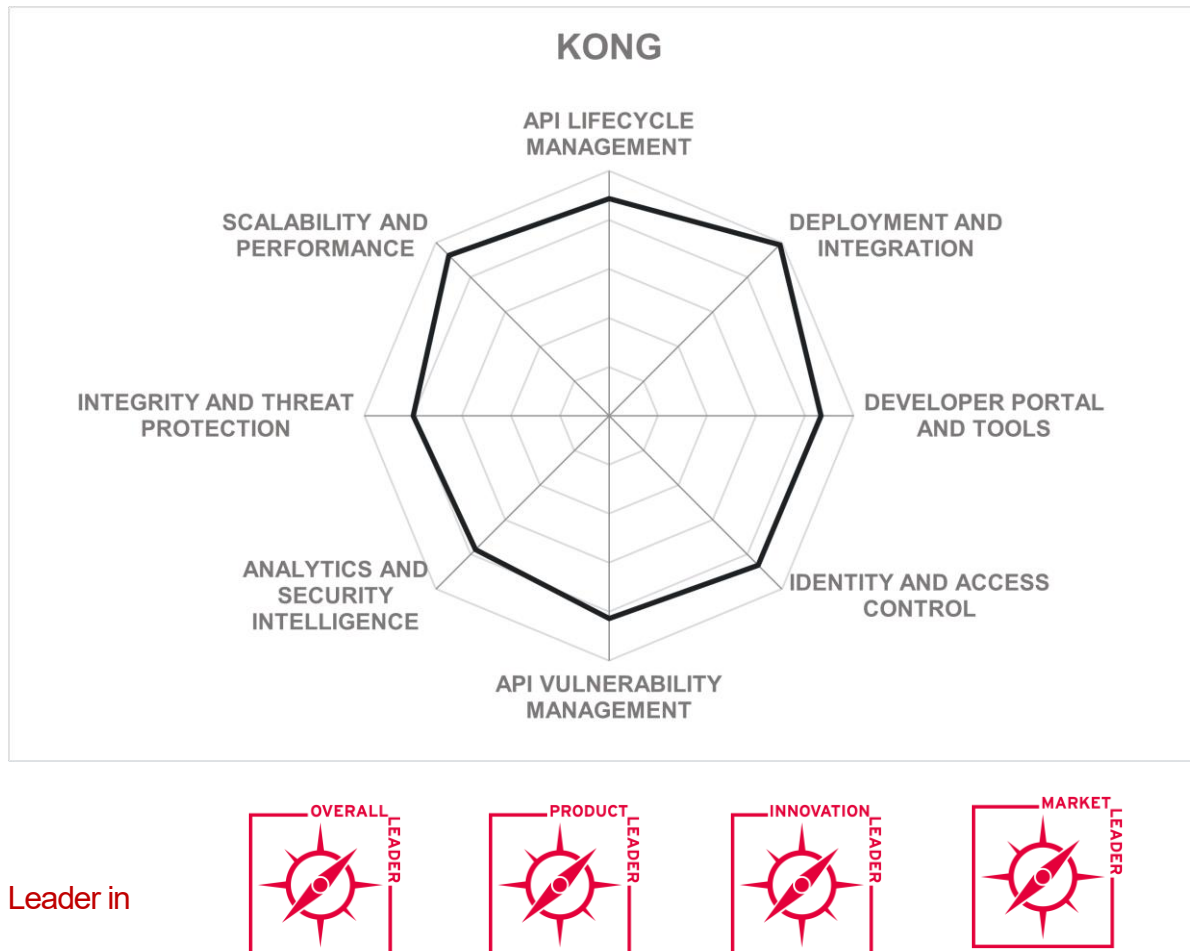
Strengths

- Unified protection for both web applications and APIs via a single platform.
- Support for REST, GraphQL, gRPC, SOAP, and JSON-RPC APIs.
- Flexible deployment across cloud WAF, gateway, and service mesh environments.
- Inline detection and response capabilities with real-time blocking.
- Integrated business logic abuse detection and policy enforcement.
- Strong analytics with Attack Analytics and ML-driven incident correlation.
- OpenAPI-based positive security model with schema validation.
- Advanced bot protection integrated natively into the platform.
- GenAI-aware threat detection and protection for AI-powered applications.
- Centralized management console with global policy control.
- Transparent, privacy-conscious data processing and residency controls.

Challenges

- No support for fine-grained access governance or policy workflows.
- GenAI-specific protections are emerging but still under development.
- Full feature parity between cloud and on-prem deployments is yet to be achieved.

Kong – API Platform



Kong Inc., founded in 2015 and headquartered in San Francisco, CA, has evolved from an open-source API gateway provider into a comprehensive platform vendor for managing APIs, services, and event-based architectures. The company has steadily grown into a global presence with a customer base that spans industries and geographies. Kong positions itself as a strategic enabler for modern enterprise architectures by offering an API platform that supports AI governance and federated API management at scale.

Its flagship offering, the Kong API Platform, provides a unified infrastructure to design, build, test, govern, and secure APIs across diverse environments. At the heart of this vision is Kong Connect, a cloud-native SaaS platform that facilitates federated governance while supporting self-service capabilities for distributed teams. Together with Kong Insomnia, the

company's API design and testing tool, the platform provides full lifecycle support for both providers and consumers of APIs.

The platform includes a broad suite of capabilities designed to serve the full API lifecycle. API design, testing, and development begin with Kong Insomnia, an environment supporting collaboration among developers, architects, and business stakeholders. Insomnia facilitates OpenAPI specification design, automated testing, policy packaging, and version management. Integration with version control and CI/CD tools allows APIs to be pushed to Git repositories, triggering automated pipelines that validate, deploy, and test API configurations, supporting scalable and compliant deployments.

Discovery and inventory management are driven by Kong's Service Catalog, which aggregates data across Kong runtimes and third-party environments. It enables organizations to detect undocumented or "shadow" APIs, correlate them with infrastructure metadata, and assess their compliance through built-in scorecards. The catalog integrates with infrastructure components such as service meshes, Kubernetes, and cloud providers, and offers a 360-degree view that encompasses ownership, documentation, cost centers, and more.

APIs discovered via the Service Catalog can be evaluated against security best practices, including OWASP Top 10. Policy automation and scorecards enable platform teams to enforce guardrails across the API lifecycle. Runtime security capabilities include support for authentication, rate limiting, bot detection, and threat protection via an extensive plugin system. Plugins, which can be developed in Lua, JavaScript, or Python, allow fine-grained control over traffic flows and can be applied globally or to specific routes.

Kong Konnect supports integration with SIEM and monitoring tools, enabling real-time visibility into API consumption and performance. Security analytics and compliance posture monitoring are also accessible via the Service Catalog, which helps teams assess risk exposure and track remediation efforts. Kong's platform offers audit trails and fine-grained visibility across the control and data planes.

Kong supports various API styles and communication patterns, including REST, gRPC, GraphQL, Kafka events, and LLM interfaces. Kong's gateways enforce transport and application-layer security policies while providing resilience against common attack vectors. Kong's AI Gateway adds another layer of AI security, offering routing, observability, and policy enforcement for LLMs and agentic systems.

Kong Konnect supports deployments across public and private clouds, on-premises, and serverless environments. It enables centralized management while allowing decentralized teams to deploy and configure local gateways that conform to global governance policies. Dedicated cloud deployments are available for customers seeking fully managed infrastructure with isolation per tenant.

The platform supports numerous third-party tools across its development, security, and observability stacks. These include service meshes, CI/CD tools, IAM platforms, and dedicated API security solutions like Traceable. Customers can embed Kong into their

existing DevOps pipelines, automate deployments using Terraform and Kubernetes operators, and leverage prebuilt plugins or build custom logic.

Kong's federated governance model allows central platform teams to retain control while empowering application teams with self-service infrastructure provisioning. The AI Gateway, which supports universal interfaces and runtime protection for AI services, is one of the first of its kind in the market. The Service Catalog's scorecard-based governance, contextual discovery, and integration with non-Kong runtimes stand out for their developer focus and enterprise readiness. However, Kong's security capabilities, while broad, are not yet as deep or specialized as those offered by pure-play API security vendors. Areas like runtime anomaly detection or data leak prevention remain potential growth opportunities.

Its platform appeals to large enterprises undergoing API-led modernization, particularly those aiming to consolidate fragmented toolchains, enforce consistent governance, and accelerate digital innovation. Organizations with complex hybrid or multi-cloud environments, as well as those investing in LLM-based solutions, will find particular value in Kong's capabilities.

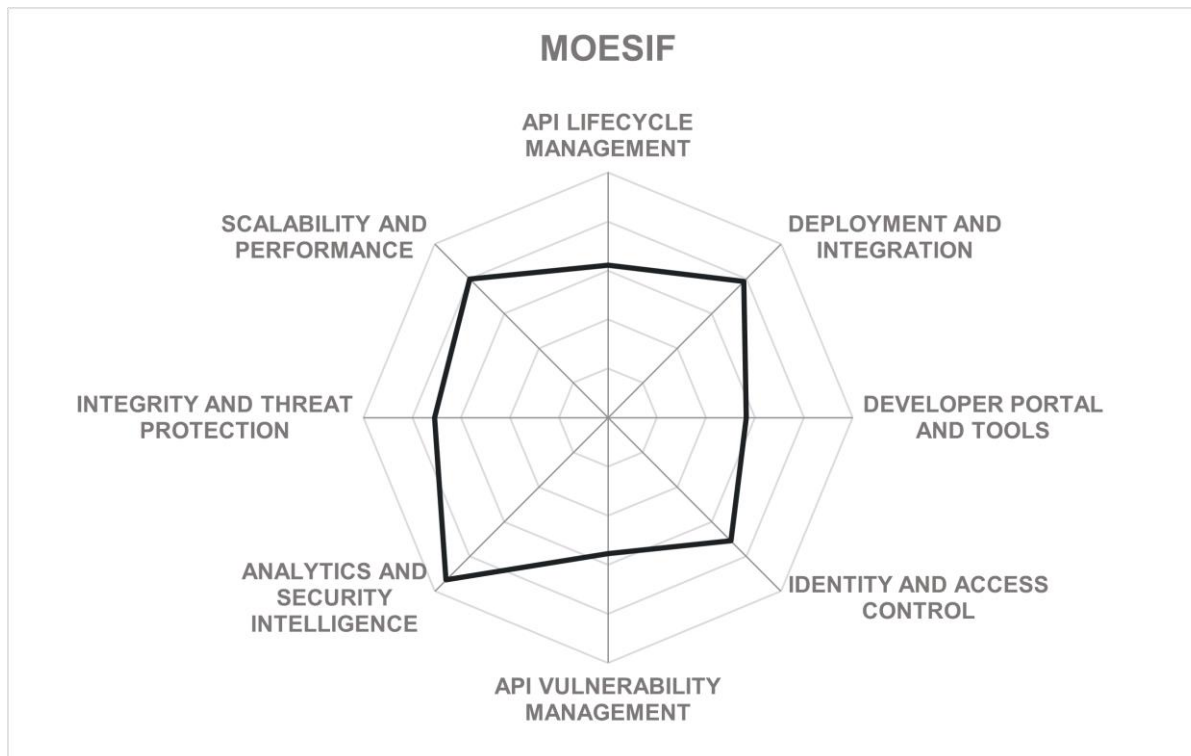
Strengths

- Strong developer tooling with Kong Insomnia and automation-first philosophy.
- Federated governance model enables scalable platform operations.
- Hybrid and multi-cloud deployment flexibility, including dedicated cloud options.
- AI Gateway for managing LLM workloads via unified policies.
- Broad plugin ecosystem with extensibility across traffic control and security.
- Advanced Service Catalog with contextual discovery and governance scorecards.
- Integration with third-party observability, and security tools.
- Deep CI/CD and IaC support, including Terraform, Kubernetes, and GitOps.
- API monetization enablement with third-party partnerships and plugins.

Challenges

- API threat detection and protection capabilities are still developing, but provides integrations with pureplay security solutions.
- Observability and anomaly detection may require third-party tools.
- Breadth of features may increase onboarding time for smaller teams or SMBs.

Moesif – Moesif



Founded in 2016 and headquartered in San Francisco, CA, Moesif is a software vendor focused on providing API analytics, observability, and monetization solutions. With a customer base that includes financial institutions, SaaS platforms, and GenAI startups, Moesif has carved out a niche among enterprises seeking a data-driven approach to managing and monetizing APIs. The company offers a SaaS platform, currently focusing its business operations primarily on the US and European markets.

Moesif positions itself not as a traditional API security platform but as a comprehensive solution for API productization. The company's vision emphasizes the convergence of observability, customer engagement, and monetization to drive better business outcomes through APIs. Moesif's product is aimed at API providers looking to maximize the business value of their interfaces by offering rich analytics, intelligent alerting, and advanced governance capabilities. The solution also integrates monetization features such as metering, usage-based billing, and quota enforcement.

The core product includes an intuitive graphical user interface for constructing dashboards, reports, and alerts. These dashboards, organized by function (e.g., engineering metrics, business performance, customer engagement), come preconfigured with each integration

and can be tailored further to suit customer-specific KPIs. Moesif supports integrations with a wide range of API gateways (such as Kong, WSO2, and Amazon API Gateway), infrastructure tools, billing providers, and support platforms. Customers can also configure data ingest via client-side SDKs or through open APIs and use the optional secure proxy mode with BYOK encryption for enhanced privacy compliance.

Moesif focuses strongly on production-stage monitoring and business analytics. This includes discovery and categorization of both internal and third-party APIs, enriched with support for OpenAPI specification validation. In addition, it supports conformance monitoring and behavioral anomaly detection through advanced machine learning, flagging suspicious activities. Privacy controls ensure differential access to sensitive data based on role-based or attribute-based access rules, with built-in GDPR support.

Moesif provides granular visibility into API traffic down to individual payloads through full-body analytics (optionally enabled). Customers can define custom metrics using a built-in scripting language and trigger alerts for various anomalies. AI-based features such as “AI Explain” allow users to query data using natural language and receive summarized insights or suggestions for further analysis. Dashboards can track business metrics such as growth rate by customer or usage by feature, as well as technical metrics like error rates or latency. Alerting and webhook-based workflows allow for integrations with external tools.

Moesif’s API governance capabilities include sophisticated quota enforcement and blocking logic based on business rules. This can be used to automate service degradation or access denial in response to abnormal patterns. The system supports prepaid and postpaid billing models with real-time consumption tracking, and allows customers to define products, pricing models, and subscription plans within Moesif itself or integrate with external billing providers.

The platform’s hybrid, multi-cloud deployment model enables visibility and policy enforcement across mixed environments. Moesif supports federated integrations across multiple gateways and cloud regions, allowing organizations to maintain consistent analytics and governance across all infrastructure. Optional client-side encryption can address data privacy concerns for regulated industries. AI security functionality extends to detecting threats against generative AI APIs, such as excessive token consumption or unauthorized access, though it does not address prompt injections.

Moesif’s key differentiators lie in its emphasis on the business side of API productization. Its unique strengths include a rich set of features around monetization alongside deep observability. The GUI-driven experience and rapid deployment further differentiate the platform from heavier, developer-centric observability tools. However, the product currently lacks coverage of early API lifecycle phases and does not offer some traditional API security features such as content filtering.

Moesif is well suited for product and engineering teams in enterprises with commercial API strategies, especially those needing actionable insights and monetization capabilities. Typical use cases include GenAI platforms with token-based pricing, financial data APIs requiring user-level metering, and SaaS platforms seeking embedded usage analytics.

In late May 2025, Moesif's acquisition by WSO2 was announced. As part of the agreement, it will operate as an independent subsidiary under WSO2's API Management Business Unit. The Moesif brand and current product offering will be retained, and its leadership along with its team will continue to drive existing business and expand customer growth globally.

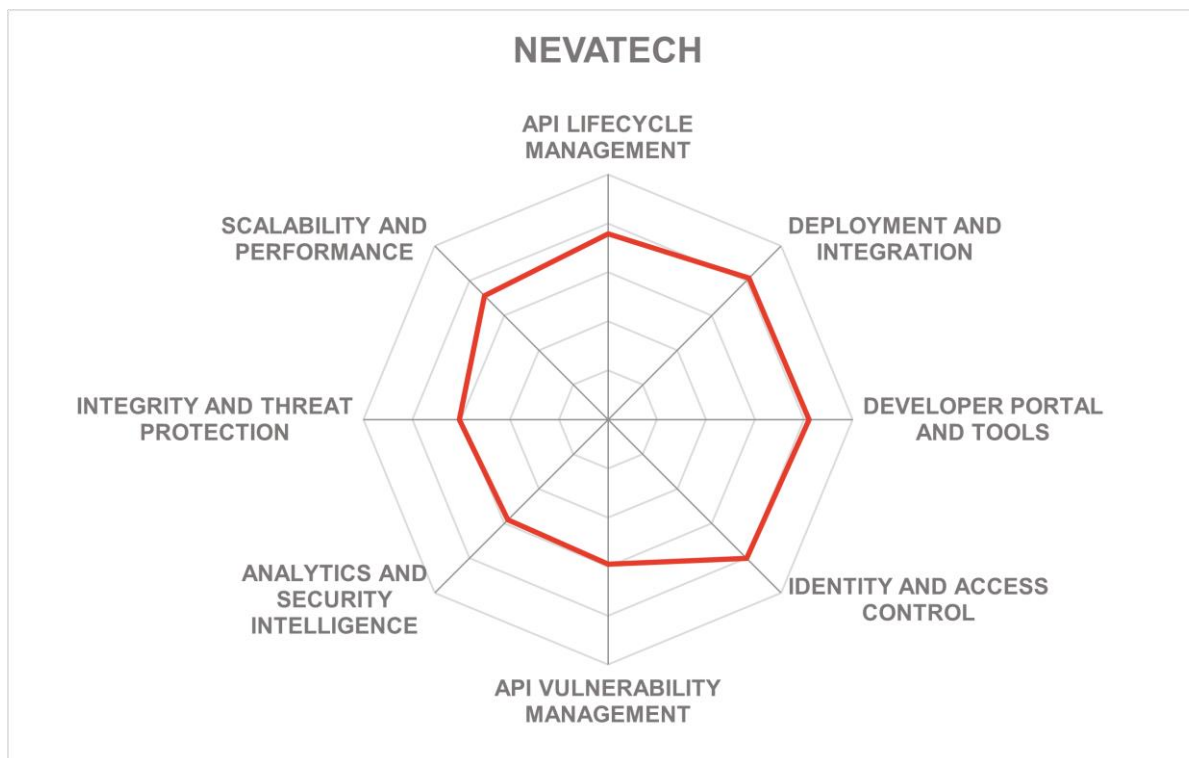
Strengths

- Comprehensive API observability with full payload analysis and customizable dashboards.
- Strong support for API monetization, including usage-based billing and pricing model management.
- Out-of-the-box integrations with popular gateways, billing, and support platforms.
- GUI-driven workflow with low learning curve and fast deployment.
- AI-powered features for report generation and anomaly detection.
- Secure proxy mode with client-side encryption and BYOK support.
- Role-based privacy controls for GDPR compliance.
- Open-source developer portal and embedded dashboards for customer visibility.

Challenges

- No built-in capabilities for vulnerability scanning or attack surface management.
- Limited support for strong authentication.
- No out-of-the-box integration with SIEM/XDR platforms.
- Cannot be deployed in a fully isolated or on-prem mode for regulated scenarios.

Nevatech – Sentinel API Management



Nevatech is a privately held software vendor founded in 2011 and headquartered in Atlanta, GA. The company is focused on delivering robust, flexible, and developer-friendly solutions for API management and governance. Over the years, Nevatech has gained recognition for its strong alignment with Microsoft-centric environments. However, with the latest evolution of its platform, the company is strategically expanding its footprint to address modern enterprise requirements for cross-platform API security, scalability, and automation across hybrid and cloud-native deployments.

Sentinet API Management is Nevatech's flagship product, designed as a platform for the full lifecycle of API management. While previous versions of Sentinet were built exclusively on the .NET Framework and required Windows Server for deployment, the current architecture is reengineered with .NET Core at its foundation. This shift enables seamless deployment on modern operating systems, including Linux, and better integration with containers and Kubernetes-based environments. The new architecture supports both backward compatibility (branded as "Sentinet Classic") and a more modular model for contemporary API management use cases.

Sentinet provides a full set of capabilities for designing, managing, securing, and monitoring APIs. The platform supports REST and SOAP APIs and integrates smoothly with existing systems with protocol transformation and legacy service enablement. APIs can be created from scratch or imported from existing services, and developers are offered a graphical designer, OpenAPI support, and tooling for testing with mock responses. The product includes a developer portal for onboarding, key provisioning, and tiered access management. The entire platform is API-first and automatable through both management APIs and a command-line interface (CLI), allowing users to fully script deployments and configurations.

Users can define new services starting from scratch or from an OpenAPI specification and incrementally build out backend services with virtual APIs using a drag-and-drop visual interface. Virtual services can be tested before connecting to physical backends using mock responses, effectively decoupling design and implementation phases. The platform also provides extensive versioning support for services and policies. Notably, policy definitions can be embedded directly into OpenAPI specifications.

Discovery and inventory management features enable users to track APIs across the organization, including external and third-party APIs. Backend APIs are managed independently of virtual (gateway-facing) APIs, which supports reuse, simplifies updates, and enables multiple gateway APIs to share a single backend. This separation significantly enhances scalability of both operations and governance.

The platform can detect a broad spectrum of policy misconfigurations and supports enforcement of OpenAPI-based constraints. Access policies are separated into authentication (policies) and authorization (access rules), encouraging users to design their API security in accordance with best practices. Each access rule can combine multiple conditions, such as claims, time of day, IP filters, rate limits, and message content, and may include user-defined logic via extensible components written in C#.

Sentinet provides a centralized management console with support for policy administration, operational monitoring, and extensive reporting. All API transactions, policy changes, and access decisions are logged with full auditability. Reports can be generated on demand or scheduled and exported for compliance or forensic analysis. The system supports real-time anomaly detection and SIEM integration via Syslog.

The platform offers multi-layered protection for API communications, including TLS and mTLS, OAuth 2.0, OpenID Connect, JWT, and SAML. A key design principle is the strict separation of policy-based security from message processing logic. Message processing workflows are designed independently, enabling complex logic like circuit breaker patterns, token exchange, or data transformation. Customers can implement non-standard identity orchestration scenarios, which is useful when integrating with legacy or non-compliant third-party APIs.

Unlike many competitors who merge policy enforcement and message transformation into monolithic definitions, Sentinet enforces a clean separation between authentication, authorization, and message processing. Its strict interpretation of “policy” as a declarative,

client-facing construct is reflected in its OpenAPI alignment and design-time validation. Extensibility is another strong point: users can embed custom components throughout the platform's runtime logic, enabling unique and complex integration scenarios. However, the user interface of the new console remains relatively unpolished, and the absence of AI-based features or built-in threat intelligence limits Sentinet's appeal to security-forward organizations.

Sentinet supports all major deployment models, including on-premises, private and public clouds, and hybrid configurations. Containerization and Kubernetes support is available out of the box, and customers can use official Docker images or build their own. This infrastructure flexibility, coupled with a fully scriptable deployment model, enables Sentinet to operate efficiently in modern DevOps and infrastructure-as-code environments.

Sentinet's customer base is traditionally skewed toward enterprises with strong Microsoft investments, but the shift to .NET Core and expanded deployment options make it increasingly attractive to a broader range of mid-size and large organizations. Organizations requiring flexible, fine-grained API governance and lifecycle management, especially those with legacy systems to modernize, should consider evaluating Sentinet.

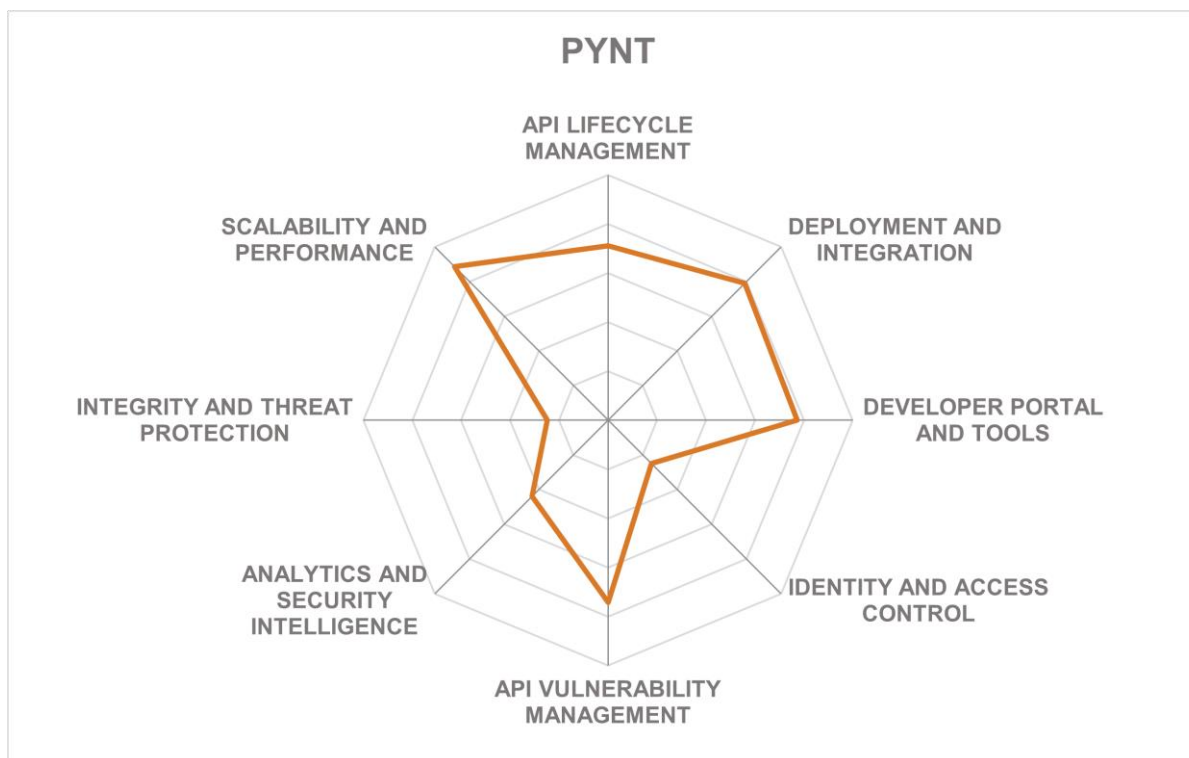
Strengths

- Strong support for Microsoft technologies and environments, including native Windows authentication.
- Clear separation of authentication, authorization, and message processing.
- Fully redesigned architecture built on .NET Core for cross-platform deployment.
- Declarative policy model compatible with OpenAPI specifications.
- Fine-grained access control with reusable access rules and condition combinations.
- Advanced extensibility via custom .NET components.
- Supports design-first API development with full mock testing capabilities.
- Independent lifecycle and inventory management of backend and virtual APIs.
- Scriptable deployment with CLI support for full automation.

Challenges

- No current support for AI/ML or generative AI features.
- No built-in support for GraphQL, gRPC, or event-based messaging protocols.
- User interface for the management console still lacks polish.
- Customization requires familiarity with .NET and C#.
- Weak market visibility and partner ecosystem relative to larger vendors.

Pynt – API Security



Founded in 2022 and headquartered in Tel Aviv, Israel, Pynt is a young cybersecurity vendor that specializes in automated, AI-enhanced API security testing. Despite its relatively short history, the company has already established a presence among security-conscious enterprise users and developers, with a rapidly growing footprint. Pynt's leadership team brings extensive experience from the automotive security domain.

Pynt's strategic vision revolves around addressing the increasingly complex and fragmented nature of application security, where the traditional boundaries between web applications, APIs, and AI-powered agents are no longer relevant. Pynt delivers a novel approach to API security by combining intelligent API discovery, contextual risk analysis, and automated testing. Integrated into CI/CD and design workflows, it empowers both developers and security teams to proactively detect and remediate vulnerabilities early, without requiring complex runtime deployments.

Pynt's core product is a fully integrated suite that delivers automated API and application security testing capabilities, operating both as a SaaS platform and a deployable container

for on-premises and hybrid use. The platform supports over 30 integrations, including popular CI/CD tools, Postman, VS Code, GitHub Actions, and proxy- or traffic-based inspection mechanisms such as eBPF. Its standout capability is its AI-powered attack simulation engine, which leverages LLMs for both test generation and contextual analysis. Pynt can test modern APIs (REST, GraphQL), LLM-based APIs, and traditional web applications, delivering results with comprehensive evidence, remediation suggestions, and auto-ticketing into systems like Jira.

The platform's API discovery engine combines traffic analysis, OpenAPI spec ingestion, and runtime observation to produce a continuously updated inventory of all internal, external, and third-party APIs. Shadow APIs, undocumented endpoints, and discrepancies between source artifacts are flagged automatically. This inventory is further enriched with metadata such as authentication status, risk score, and sensitive data exposure. Findings are linked to real requests and responses and can be reproduced via automatically generated cURL commands.

Rather than relying solely on static analysis or passive monitoring, Pynt simulates real attacks by launching active probes against APIs. These cover a wide range of vulnerabilities, from OWASP Top 10 to LLM-specific threats like prompt injection. The system dynamically adjusts test scenarios based on contextual factors such as user roles, parameter semantics, and endpoint function. This approach allows Pynt to detect complex business logic flaws, including BOLA, excessive data exposure, and improper session handling. For each vulnerability, Pynt presents detailed findings including CWE classifications, application flow traces, and supporting evidence. A penetration test report generator supports compliance needs with standardized documentation. An AI insights module offers more in-depth analysis, including regression detection and effectiveness of remediation over time.

Pynt delivers rich remediation workflows: AI-enriched vulnerability descriptions, full request/response logs, evidence and repro scripts, plus false-positive handling and ticketing integrations. This ensures fast triage and resolution across security and engineering. The vulnerability lifecycle can also be automated as part of the pipeline. However, while the platform provides extensive vulnerability data and reporting, it does not currently offer runtime protection or threat mitigation mechanisms, and does not yet map findings to MITRE ATT&CK.

Pynt's dual-layer AI architecture uses generative models not just to simulate attacks but to analyze responses, significantly reducing false positives. This approach is complemented by synthetic traffic generation from API specs when real traffic is sparse. Pynt can also scan for LLM-related security issues, such as jailbreak attempts and insecure chaining. Additionally, Pynt's strong partnership with Postman enables it to support developer-centric security workflows.

While Pynt provides strong pre-production testing, it lacks runtime enforcement features such as inline blocking or anomaly detection. Features like OpenAPI validation, JSON schema validation, and full policy lifecycle management are on the roadmap but not yet available. Current deployment options are skewed toward public cloud and DevOps-friendly environments, with more limited support for legacy IT or fully managed services.

Pynt's customer base spans North America, EMEA, and parts of APAC, with the strongest adoption in security-sensitive industries such as finance, aerospace, and healthcare. It is particularly well-suited for organizations with mature CI/CD pipelines and a need for proactive API security testing integrated into development workflows. Customers seeking protection for LLM-enabled applications or looking to streamline penetration testing processes will also find value in the platform.

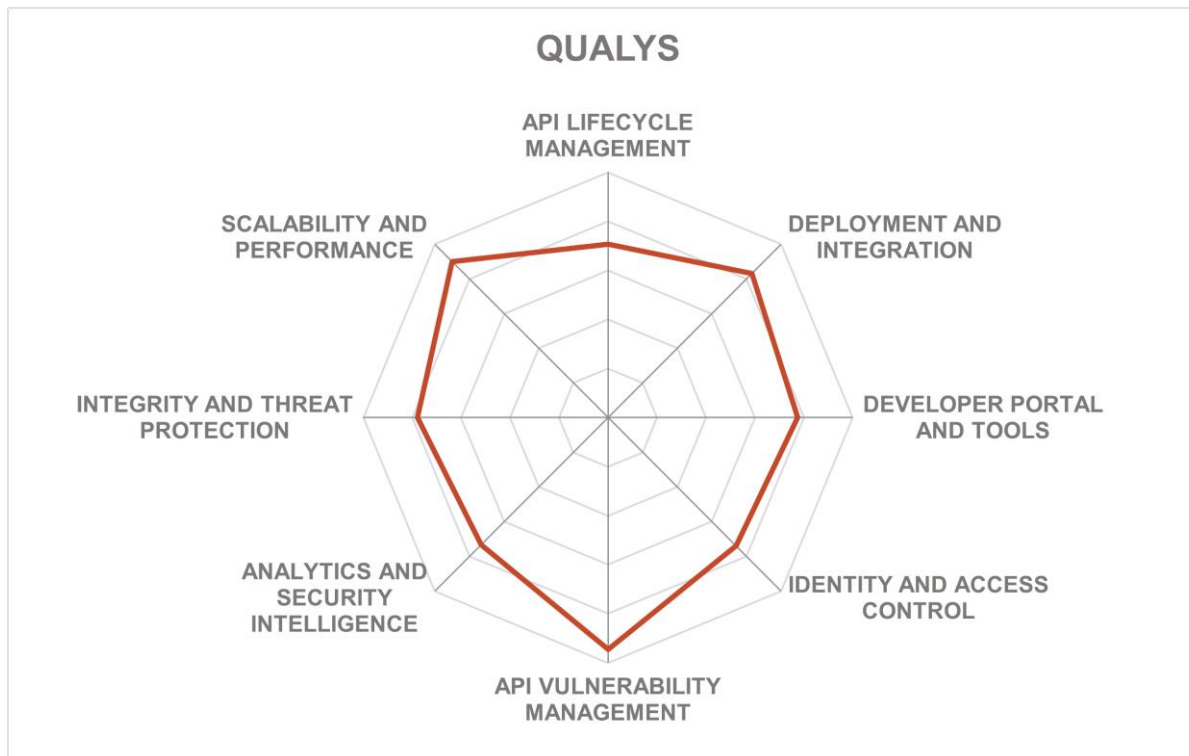
Strengths

- Comprehensive API discovery across multiple sources.
- Requires no configuration, rules, or scripts.
- Context-aware attack engine with support for business logic flaws.
- Dual-layer LLM-driven testing and response analysis to minimize false positives.
- Continuous scanning with CI/CD and IDE integrations.
- Full-stack testing including Web, API, and AI surfaces.
- Evidence-rich vulnerability reporting with ticketing and remediation guidance.
- Shadow and orphan API detection with traffic/source correlation.
- Pentest report generator supports compliance and audit requirements.
- AI-generated synthetic traffic for testing environments with limited visibility.
- Early support for LLM-specific attack vectors and AI security use cases.
- Lightweight, fast, and flexible, scales to thousands of APIs per test.

Challenges

- Lacks runtime protection or active threat mitigation capabilities.
- Limited support for legacy or event-based protocols (e.g., SOAP, Kafka).
- Minimal current integration with security data lakes, SIEM, or XDR.
- LLM testing features are innovative but still evolving and potentially fragile under non-deterministic behavior.

Qualys – TotalAppSec



Leader in



Founded in 1999 and headquartered in Foster City, CA, Qualys is a well-established cybersecurity vendor with a long track record in vulnerability management. Over the years, the company has significantly expanded its capabilities into various domains of enterprise risk management, including application security. Its TotalAppSec product represents a comprehensive extension of this portfolio into web application and API security with a unified solution for securing modern applications across on-premises, cloud-native, and hybrid infrastructures.

Qualys approaches the API security challenge through the lens of holistic risk reduction. Rather than offering a standalone API security point solution, the company has integrated API discovery and testing into its broader application security platform, which also includes dynamic web application scanning and web malware detection. This strategy is tightly aligned with Qualys' overall vision of delivering a unified platform for enterprise cyber risk

management through its TruRisk™ scoring engine, centralized policy enforcement, and extensive automation capabilities.

TotalAppSec combines three core functions in a single license and user interface: web application security, API security, and malware detection powered by deep learning. The platform automates the discovery of both web apps and APIs, including undocumented, shadow, and deprecated endpoints, using multiple approaches such as cloud gateway integrations (AWS, Azure, MuleSoft, Apigee), log parsing, static OpenAPI spec analysis, and even scanning external internet exposure. Qualys provides dynamic security assessment against OWASP Top 10 risks, detection of sensitive data exposures, and OpenAPI conformance validation with over 200 test signatures.

Vulnerability detection and prioritization are deeply integrated with TruRisk, which combines asset criticality, exploit context, and threat intelligence from over 25 feeds to rank API and web app risks. These insights power remediation workflows that are enriched by integrations with ITSM systems like ServiceNow and JIRA, as well as CI/CD platforms such as Jenkins, Azure DevOps, and GitHub Actions.

The scope of API security within TotalAppSec includes discovery and classification of all APIs across environments, risk-based assessment based on usage and exposure, and support for REST, SOAP, GraphQL, and gRPC protocols. The product validates API implementations specifications and identifies OWASP API Top 10 violations, including lack of rate limiting, improper authentication, and sensitive data exposure. Qualys complements its detection capabilities with AI-powered threat analysis, clustering vulnerabilities to reduce scan time and increase accuracy. The company claims improvements of up to 80% by focusing on high-risk patterns identified through historical analysis.

TotalAppSec is tightly integrated with Qualys' newly introduced AI Security Posture Management module. This allows customers to assess risks related to AI models, workloads, and APIs interacting with LLMs. Emerging threats such as prompt injections and model jailbreaks are addressed using proprietary detection logic.

A key differentiator for Qualys is the convergence of app and API security with other layers of enterprise risk visibility. By consolidating vulnerabilities from third-party pen tests, static code analyzers, and bug bounty platforms, it provides centralized insight and orchestration across the full application stack.

That said, there are tradeoffs to Qualys' unified approach. The platform does not currently offer in-line API gateway or real-time request blocking capabilities, nor does it provide out-of-the-box runtime protection against volumetric denial-of-service attacks. Instead, its protection is primarily derived from vulnerability prevention and threat-informed scanning.

Qualys's licensing model reflects its platform-first approach: a single SKU covers both web apps and APIs, allowing customers to dynamically reallocate licenses based on changing needs. This model has proven attractive to large customers transitioning from siloed API security tools or traditional DAST vendors.

Qualys has a well-diversified global customer base, including major financial institutions, governments, and regulated enterprises across all geographic regions. The solution is especially relevant for organizations seeking to consolidate vulnerability and API security into a central platform, those looking to modernize legacy DAST workflows, and enterprises prioritizing risk-based remediation and governance across web and cloud application environments.

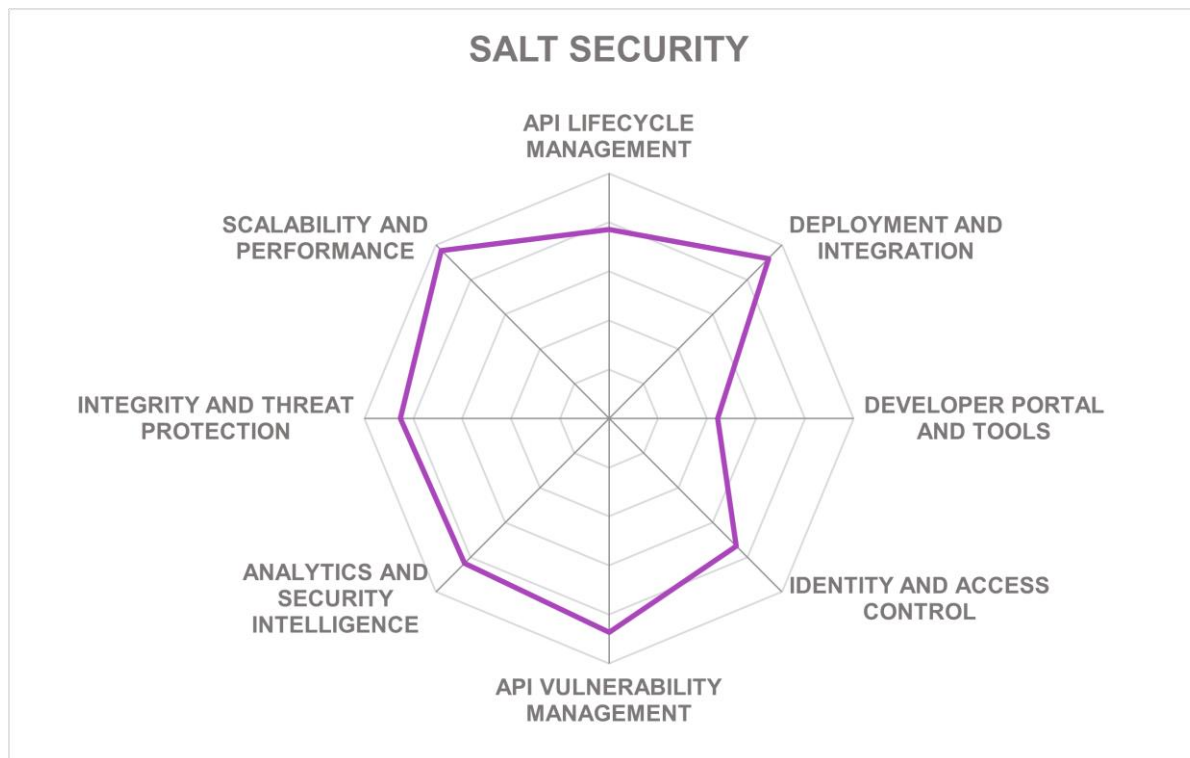
Strengths:

- Unified platform for web apps, APIs, and malware detection with shared workflows.
- Extensive API discovery capabilities covering gateways, cloud, source code, and traffic analysis.
- Deep integration with CI/CD pipelines and ITSM tools.
- TruRisk™ scoring provides risk-based prioritization across applications and infrastructure.
- Comprehensive OWASP Top 10 testing for both web and API vulnerabilities.
- AI-powered Quick Scan reduces scan times significantly while maintaining high accuracy.
- Integration with external bug bounty and SAST tools for consolidated findings.
- Native support for deep learning-based malware detection in web assets.
- Innovative licensing model allows flexible allocation between web app and API coverage.

Challenges:

- Lacks in-line API request protection or API gateway-like enforcement capabilities.
- Does not currently provide real-time mitigation features such as automated rate limiting or anomaly blocking.
- Certain advanced runtime behaviors (e.g., GraphQL introspection abuse) may require manual detection tuning.
- Some AI security capabilities are still in early stages of maturity.

Salt Security – API Security Platform



Salt Security, founded in 2016 and headquartered in Palo Alto, CA, is a privately held cybersecurity vendor specializing in API security. With its origins rooted in the Israeli cybersecurity community, Salt has established itself as one of the early and prominent players in the API security market. The company's strategic direction is delivering comprehensive, enterprise-grade API threat protection, discovery, and governance capabilities through a data-driven AI-powered platform. Salt positions itself as a non-intrusive yet deeply integrated solution that complements existing security infrastructure.

Rather than replacing API gateways, WAFs, or CSPMs, Salt augments them by addressing their well-known visibility and protection gaps. Salt's platform emphasizes reducing risk before attacks occur and detecting sophisticated threats as they unfold. It targets large enterprises with complex, dynamic application environments. The company has invested heavily in a growing ecosystem of technical partners and tight integrations with vendors such as CrowdStrike, enabling broader coverage and ease of deployment.

The Salt API Security Platform is delivered as a suite of integrated capabilities that include agentless, out-of-band API discovery, policy-driven posture governance, and AI-powered threat detection. The platform can be deployed as a cloud-native solution or in a hybrid mode, where sensitive data never leaves the customer's environment. Its architecture supports integration with existing traffic sources including API gateways, CDNs, load balancers, and Kubernetes clusters, offering real-time insights without performance impact. Salt's telemetry is processed in a cloud-based data lake, enabling analysis across vast API datasets to support behavioral baselining, risk scoring, and threat modeling.

Salt provides automated continuous API discovery across multi-cloud, container, and serverless environments. The Cloud Connect integration, particularly with AWS, Azure, and GCP, allows for instant asset inventory without the need for sensor deployment, while the CrowdStrike Foundry integration provides simple sensor deployment throughout a CrowdStrike environment. The platform distinguishes between critical business APIs and low-risk or internal ones, avoiding alert fatigue and promoting operational focus. Discovered APIs are evaluated against OpenAPI specs and posture rules, helping to identify undocumented or misconfigured interfaces, including shadow APIs behind internal gateways.

Salt's posture governance engine, where users can apply predefined or custom-built rules, identifies gaps such as missing authentication, excessive data exposure, or non-compliance with internal security standards. These posture gaps are not static; they reflect live traffic observations and configuration data. Policy Hub offers more than 70 out-of-the-box rules that map to frameworks such as PCI DSS, HIPAA, and NIST. Remediation guidance is embedded directly into the UI, and integration with ticketing tools like Jira or ServiceNow supports developer engagement.

At runtime, Salt offers robust analytics and threat detection using a patented AI/ML engine. This engine differentiates between benign anomalies and actual attacks with a claimed 92% intent accuracy. It performs real-time attack correlation, displaying actionable attack timelines, source data, and natural-language summaries for security analysts. The platform detects advanced attacks like low-and-slow exfiltration, SSRF, and abuse of business logic vulnerabilities. Events are enriched with MITRE ATT&CK mappings and can be forwarded directly to SIEM and XDR platforms, including native integrations with CrowdStrike.

The platform's threat protection mechanisms extend to secure data transmission, detection of payload tampering, and identification of common and API-specific attack vectors. Salt supports all major API protocols including REST, GraphQL, SOAP, and gRPC, and is capable of detecting threats across them. While Salt does not perform inline blocking itself, it integrates with WAFs, firewalls, and gateways for policy enforcement and mitigation. Advanced automation capabilities are available through the Workflow Hub, allowing users to build custom playbooks for threat response. Salt has also made early moves into AI security by offering guardrails against prompt injection and other LLM-specific threats.

Customers can choose fully cloud-based deployment or hybrid models to align with their privacy and compliance requirements. Salt's architecture is environment-agnostic and scales across multi-cloud, Kubernetes, and microservice landscapes. Integration with CrowdStrike's

Foundry enables sensor deployment via endpoint agents, streamlining rollout in complex enterprise environments. The platform's focus on metadata processing allows enterprises to maintain strict data residency and sovereignty, while still benefiting from centralized analytics and detection.

Key differentiators for Salt Security include its posture governance engine, which acts as a repository for policy enforcement, and its patented intent analysis engine that dramatically reduces alert volumes while preserving detection accuracy. The deep integration with CrowdStrike Foundry and NG-SIEM allows for rapid sensor deployment and real-time telemetry correlation. Another strength is the platform's ability to function in a "headless" mode, embedding seamlessly into customer environments without requiring adoption of its own interface. However, the lack of native enforcement and limited developer capabilities may limit appeal for organizations seeking an all-in-one API security solution.

Salt primarily targets large enterprises in regulated and data-sensitive industries such as finance, retail, manufacturing, and aviation. The solution is well-suited for organizations with mature security programs that need to close gaps in API visibility and threat protection without disrupting existing architectures. Customers operating hybrid or multi-cloud environments, especially those already invested in CrowdStrike or modern SIEMs, will find strong alignment with Salt's integration strategy.

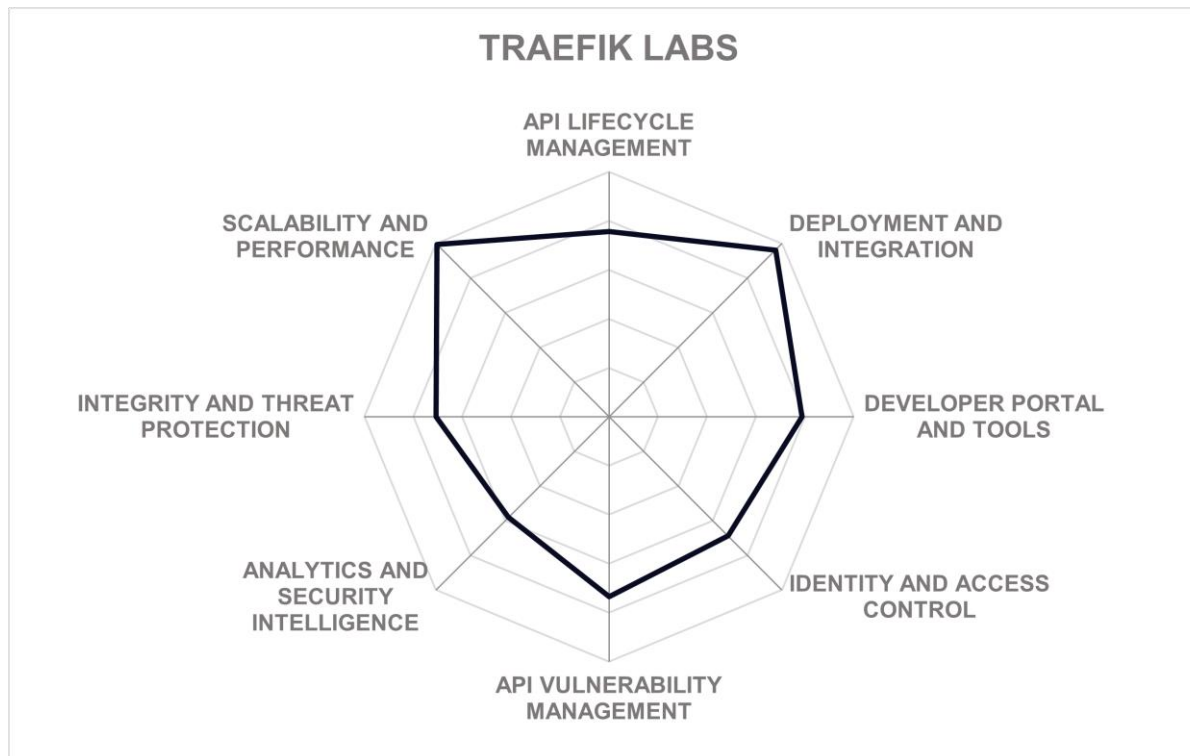
Strengths

- Out-of-band architecture ensures no performance impact and simple deployment.
- Support for a wide range of API protocols and runtime environments.
- CrowdStrike integration enables frictionless sensor deployment at scale.
- Cloud Connect provides near-instant cloud API discovery without deploying agents.
- Unified inventory with detailed risk scoring, metadata, and posture visibility.
- Patented intent analysis engine minimizes false positives with high threat detection accuracy.
- Deep integration with CrowdStrike NG-SIEM and firewall for actionable analytics and blocking.
- Posture Governance engine with prebuilt and custom compliance policies.
- Extensive automation and workflow capabilities for response orchestration.

Challenges

- No native enforcement or inline blocking capability.
- Minimal focus on developer tools and onboarding experience.
- Lacks integrated API rate limiting and SLA management (relies on gateways).
- Data privacy controls rely on metadata abstraction, which may limit use cases with stricter sovereignty requirements.

Traefik Labs – Traefik Hub Platform



Leader in



Traefik Labs, founded in 2016 and headquartered in Lyon, France, is a technology vendor with deep roots in open-source software and a strong focus on cloud-native infrastructure. The company is best known for its flagship application proxy product but also offers an open-source API gateway and its commercial Traefik Hub Platform, which together provide a complete runtime stack for discovering, routing, securing, and governing APIs. With a vision grounded in Kubernetes-native architecture, declarative configuration, and GitOps principles, Traefik aims to streamline and secure the deployment and lifecycle management of all APIs.

Rather than attempting to cover the entire API lifecycle, the company has opted to focus exclusively on runtime functionality, delivering API management, observability, and governance at scale. Their platform is deliberately designed to coexist with third-party tools, promoting interoperability with existing security and development environments. This includes support for OpenTelemetry, integration into CI/CD pipelines, and the ability to plug

into external WAFs or access governance solutions. Traefik's emphasis on a fully declarative configuration model enables traceability, policy automation, and seamless integration into modern DevSecOps workflows.

The Traefik Hub Platform comprises a suite of tools designed for progressive adoption: starting with the open-source API Gateway, advancing to the commercial API Gateway with extended security features such as WAF and AI Gateway, and culminating in full API Management with capabilities like API mocking and runtime governance. Migration across these tiers is smooth, requiring only a binary swap with configuration persistence. The platform supports a range of protocols including REST and gRPC and can be deployed in containers, virtual machines, or on bare metal. Support for the new Kubernetes Gateway API and backward compatibility with legacy ingress controllers enables gradual transitions for existing projects.

APIs can be organized into bundles, versioned, and published through a customizable developer portal. Access control, rate limiting, and quota enforcement policies are managed as code and integrated into CI/CD pipelines. Subscriptions and monetization capabilities support exposing APIs as commercial products. Mocking is fully supported for frontend/backend decoupling, allowing developers to simulate API responses during early-stage development.

Both first- and third-party APIs can be inventoried using Kubernetes integrations or traffic analysis and exposed through the Hub. API vulnerabilities are mitigated through a combination of OpenAPI contract validation, static configuration analysis, and the use of Coraza WAF for dynamic request inspection. The solution supports all OWASP API Top 10 threats and includes content filtering, request method restrictions, payload validation, and advanced threat response mechanisms. DoS/DDoS protection is supported through native capabilities and integrations with providers like Cloudflare.

Analytics and monitoring are enabled through OpenTelemetry integration, with metrics, logs, and traces available for ingestion into any SIEM or observability platform. Pre-built dashboards and support for policy compliance and SLA conformance reports are provided via tools like Grafana. The platform's security intelligence features include anomaly detection, real-time policy violation alerts, and forensic auditing of administrative activity. However, mapping to threat frameworks like MITRE ATT&CK must be implemented externally.

The AI Gateway is a recent addition designed to secure and manage APIs from LLMs and other AI models. This includes semantic caching using multiple vector database backends, content filtering, and governance for cost control and data privacy. The gateway also supports version decoupling between APIs and model backends, multi-tenant usage, and cache poisoning protection. These capabilities address growing concerns around LLM security, prompt injections, and data leakage.

The principal differentiators of the Traefik platform include its Kubernetes architecture, ease of use, developer-first design, and GitOps-native policy management. While the platform

delivers strong runtime capabilities, it currently lacks a few advanced API design-time governance and business-centric policy features found in larger competitors.

Traefik Labs serves a global customer base, spanning industries from finance and healthcare to online retail and software development. The product is particularly well-suited for organizations modernizing legacy architectures or embracing hybrid and multi-cloud deployments. Companies exploring secure API exposure for AI models will find the AI Gateway component especially valuable.

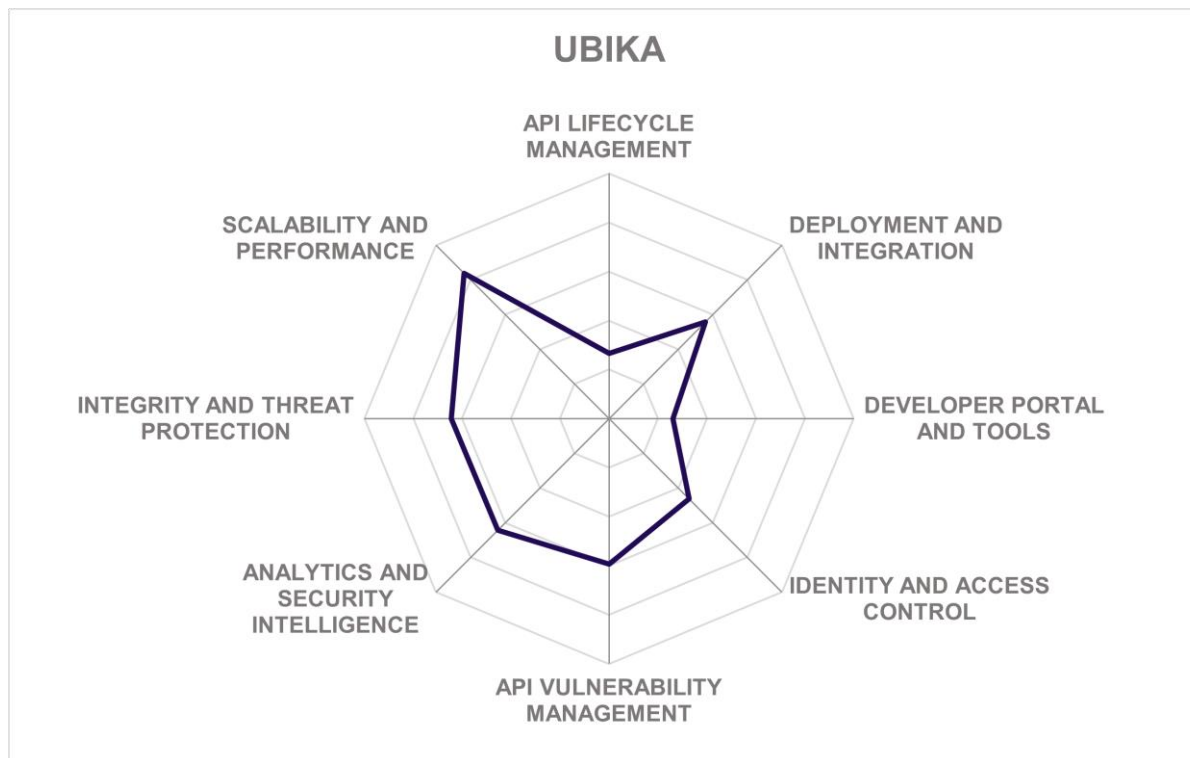
Strengths

- Kubernetes-native architecture with full Gateway API support.
- Strong protocol and deployment platform support.
- Modular, progressive adoption path from open source to enterprise API Management.
- Declarative, GitOps-native configuration for policies and deployments.
- Support for multiple vector databases and multi-tenant caching.
- Developer-friendly portal with customizable UI and API monetization.
- Semantic caching and content filtering for AI inference APIs.
- Seamless observability via OpenTelemetry integration.
- Predictable, instance-based pricing model suitable for scale.

Challenges

- Limited capabilities for early-stage API design and transformation.
- Compliance reporting is implemented by integrating third-party tools into the product.
- No support for event-driven protocols like MQTT or Kafka.
- Support for GraphQL is rudimentary at present time.

UBIKA – UBIKA WAAP Gateway and UBIKA WAAP Cloud



UBIKA is a French cybersecurity vendor founded in 2001 by the launch of DenyAll, a spin-off from Société Générale, France's largest bank. Headquartered in Meudon, France, the company has long been recognized as a European pioneer in the Web Application Firewall (WAF) space. After being a part of Rohde & Schwarz for several years, UBIKA again operates as an independent company since 2022, focusing on securing digital communications across web and API channels. Its flagship offering, UBIKA WAAP Gateway and UBIKA WAAP Cloud, exemplifies a modern interpretation of Web Application and API Protection (WAAP), integrating legacy WAF features with the latest functions aimed at holistic API security.

UBIKA positions API security as a tightly integrated but modular component of its broader WAAP architecture. This layered security approach reflects UBIKA's commitment to balancing usability with strong protection, offering customers visual, programmable interfaces for building robust policies. UBIKA emphasizes API security as a natural extension of its application protection platform. The company supports both perpetual licensing and subscription-based managed services, ensuring flexibility for customers preferring CAPEX or OPEX models.

UBIKA WAAP Gateway and UBIKA WAAP Cloud is a unified product available in multiple deployment models, including on-premises, managed services, and SaaS through European and American public cloud platforms. It supports both traditional web application security and a wide array of API protection mechanisms. The capabilities include OpenAPI 3, FHIR-compliant JSON, and XML schema validation, message signature and encryption, bot mitigation, rate limiting, anomaly detection, and an optional AI-powered antimalware engine for inspecting API payloads. The inclusion of IP reputation filtering and behavioral profiling further reinforces runtime protection.

Advanced policy workflows can dynamically adapt to conditions such as detected anomalies or threshold violations, triggering alternate enforcement profiles. This includes “under attack” policies, which can be activated manually or via SOC integrations, allowing for real-time escalation. The platform’s visual policy editor enables administrators to construct security workflows as drag-and-drop sequences, which can later be automated through Infrastructure-as-Code tools like Terraform or Ansible.

API discovery and inventory capabilities are supported through structural analysis of traffic and OpenAPI specifications. This enables automated classification of APIs and facilitates whitelisting strategies aligned with contract-first development practices. Vulnerability management is achieved by comparing live traffic against API specifications, flagging deviations that may indicate misconfigurations or attacks. Administrators can also detect excessive reliance on outdated or overly broad exception rules using the integrated Security Exception Manager. This feature helps clean up configuration drift and prevent inadvertent security gaps caused by accumulated false positives.

Real-time and scheduled reports are available to track API usage, performance, policy violations, and SLA compliance. Data can be exported to external SIEMs, allowing centralized monitoring and incident correlation. Additionally, API transactions are mapped into logical workflows, making debugging and optimization more intuitive. Monitoring and log aggregation capabilities are being enhanced as part of the roadmap, with a planned focus on simplifying incident response and improving SOC integrations.

The platform integrates with the French antimalware engine Glimps to analyze files transmitted via API calls. Administrators can define flexible thresholds, timeouts, and fallback behaviors to balance security and performance. This is particularly relevant for industries such as finance and insurance, where file uploads via APIs are common. In addition, UBIKA is experimenting with in-house machine learning engines to add risk scoring and proactive anomaly detection to their protection stack. UBIKA’s AI capabilities extend into threat detection, malware scanning, and behavioral analytics.

The company’s main differentiators include its mature, sovereign WAF heritage; its visual policy editor; and its multi-layered security model that blends positive and negative enforcement strategies. UBIKA is also the only WAAP vendor certified by ANSSI, France’s national cybersecurity agency, which adds credibility in regulated sectors. However, UBIKA’s relatively low market visibility outside of France, lack of support for some modern interfaces like GraphQL, and limited ecosystem of developer tools could impact adoption in broader enterprise or cloud-native environments.

UBIKA's customer base is predominantly located in the EMEA region, with a strong focus on French enterprises and public institutions. The platform is well-suited for customers requiring strict regulatory compliance, support for sovereign cloud environments, and in-depth customization of security policies. Organizations seeking to integrate API security into their broader web application security program, rather than deploying a separate solution, will find UBIKA a pragmatic and flexible offering.

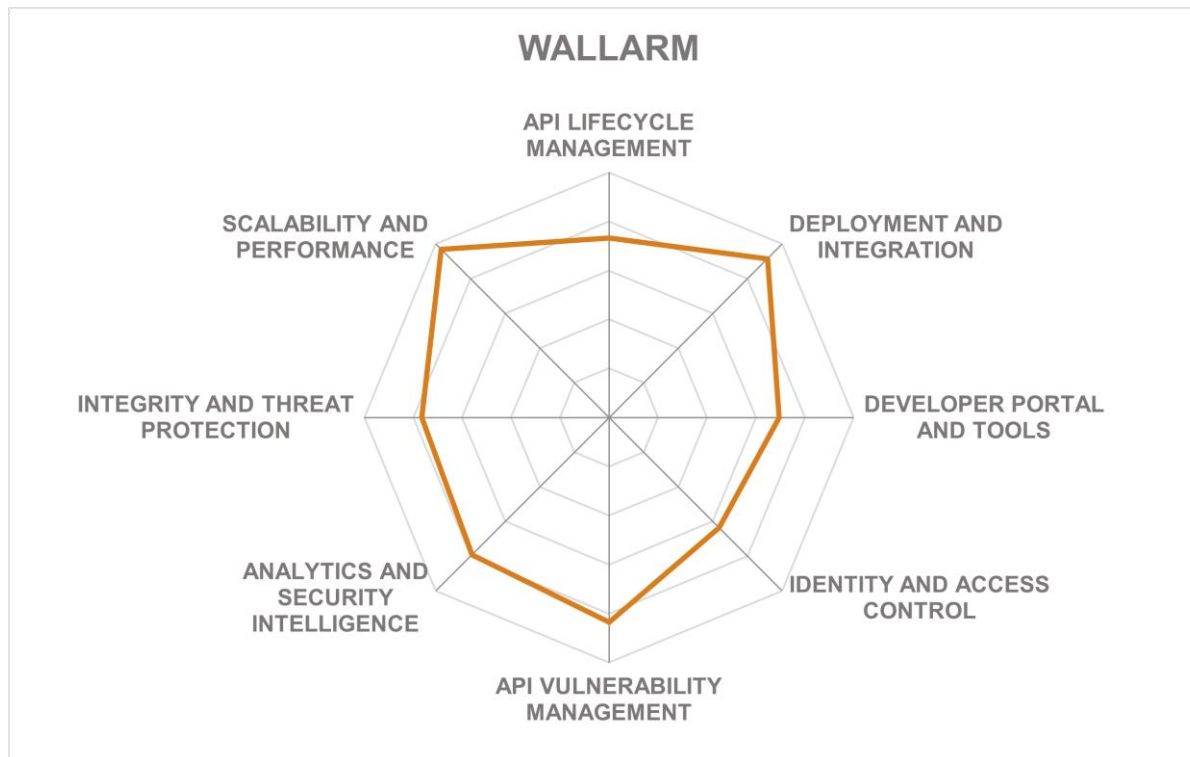
Strengths

- Long-standing expertise in web application security.
- Integrated WAAP platform with modular API security features.
- Support for OpenAPI 3.1, JSON, XML, and HL7 FHIR.
- Visual policy editor with drag-and-drop workflow design.
- Full automation support via APIs, Terraform, Ansible.
- Real-time and historical monitoring, with SIEM integration.
- Security Exception Manager for configuration hygiene.
- AI-powered malware scanning with flexible performance settings.
- Bot mitigation with good bot database and behavioral profiling.
- Certified by ANSSI (CSPN) for sovereign security compliance.
- Flexible deployment options including SaaS and managed services.

Challenges

- Limited presence outside of France and low brand visibility globally.
- Some modern API types like GraphQL are not yet supported.
- Legacy Java-based user interface pending full replacement.
- Modest ecosystem of third-party developer tools and SDKs.

Wallarm – Advanced API Security



Wallarm, founded in 2016 and headquartered in San Francisco, CA, is a cybersecurity vendor with a focus on protecting modern application architectures through a unified API security platform. The company's roots in ethical hacking and its active engagement in open-source security initiatives underscore its commitment to innovation. Over the years, Wallarm has established itself as a key player in the API security space, particularly by championing the need for real-time protection across complex and distributed environments. Its strategy reflects a strong emphasis on cloud-native scalability, deployment flexibility, and integrated machine learning to automate threat detection and response.

Wallarm Advanced API Security provides the means for identifying, assessing, and mitigating risks across the entire API lifecycle. This product is integrated into a broader Web Application and API Protection (WAAP) platform, which combines traffic monitoring, policy

enforcement, security testing, and dynamic discovery to defend against both known and emerging threats. Designed for hybrid and multi-cloud environments, it supports a wide range of protocols and standards, including REST, SOAP, gRPC, GraphQL, and WebSocket. Customers can deploy the solution inline for real-time traffic inspection or out-of-band in passive modes, achieving flexibility without compromising performance.

Wallarm uses both passive traffic analysis and active scanning to maintain a continuously updated inventory of APIs, including external, internal, and shadow endpoints. API Attack Surface Management provides visibility into usage patterns, risk posture, and sensitive data exposure. The platform highlights business-critical flows and tags APIs involved in sensitive operations like authentication or billing. Uploading OpenAPI specifications enables spec enforcement, automatically flagging and blocking endpoints that deviate from schemas.

Active vulnerability scanning, passive detection based on runtime traffic, and the Threat Replay Testing capability allow Wallarm to simulate known exploits against test or staging environments. This approach leverages real attack data to evaluate defenses in a controlled setting, delivering a higher degree of relevance than synthetic tests. Wallarm also supports enforcement of security controls such as rate limiting, session validation, and credential reuse detection. A recently added API Test Patrol capability allows users to upload an OpenAPI Specification and produce a docker container that can be integrated into their CI/CD pipeline to test APIs for vulnerabilities.

All requests and sessions are correlated to provide visibility into attack paths and threat progression. Users can explore traffic over time, analyze behaviors across sessions, and trace incidents back to their origin. These insights are supported by integrations with major SIEM and SOAR platforms like Splunk, Sentinel, and QRadar. For organizations seeking additional operational resilience, Wallarm offers co-managed SOC services that augment internal teams with expert support and policy tuning.

Wallarm's filtering nodes operate inline to detect and block malicious activity with sub-millisecond latency. This enables instantaneous prevention of attacks such as injection attempts or application-level denial of service. Protection is protocol-agnostic and includes deep inspection for GraphQL, WebSocket, and gRPC traffic. The system also identifies API abuse patterns over time, such as scraping, account takeover, and anomalous session behavior.

AI and machine learning capabilities are used to identify patterns across traffic, assess anomalies, and even assist in policy creation. The platform includes emerging protection mechanisms for AI and LLM interfaces, guarding against threats such as prompt injection and AI model abuse. The vendor treats AI security as a subset of API security, reinforcing the view that interfaces to AI systems represent an expansion of the existing attack surface.

Wallarm supports agentless, fully managed cloud deployments, on-premises installations, and hybrid models. Customers can deploy filtering nodes as container sidecars in Kubernetes clusters, as plugins for API gateways like Kong or MuleSoft. This addresses privacy-sensitive customers, regional data residency requirements, and performance optimization needs. Wallarm's Security Edge allows customers to deploy the platform simply

by redirecting DNS for their APIs. Wallarm hosts and manages the filtering nodes in public cloud providers, placing the filtering function as close to the API as possible. This provides not only easy deployment, but low latency.

Among Wallarm's differentiators are its Agentic AI protection features, its Security Edge architecture for SaaS-based enforcement, and comprehensive GraphQL policy enforcement. However, areas like code scanning remain outside of the product's current scope. Compliance automation and regulatory mapping are also limited compared to some competitors.

Wallarm's customer base is concentrated in North America, with growing presence in EMEA and LATAM. The solution is particularly appealing to mid-market and enterprise organizations operating complex, distributed environments. Wallarm has strong integration capabilities, including with CI/CD pipelines and infrastructure-as-code tools, making it suitable for security-conscious DevOps teams.

Strengths

- Unified WAAP and API security platform with integrated capabilities.
- Real-time blocking of threats inline, not just alerting.
- Broad API protocol support (REST, GraphQL, SOAP, gRPC, WebSocket) with deep protocol-specific analysis.
- Passive and active vulnerability scanning, with replay testing.
- Deep session visibility and forensic analysis tools.
- Broad spectrum of supported deployments, for inline and out-of-band scenarios.
- Extensive third-party integrations with SIEM, SOAR, CI/CD, and API gateways.
- Security Edge for easy deployment and lowest latency requirements.
- AI and Agentic AI protection features for emerging API use cases.
- Continuous update model with a co-managed SOC option.

Challenges

- No code scanning or API design validation before implementation.
- Limited native compliance automation or regulatory framework mapping.
- No support for event-based API protocols (MQTT, Kafka).
- No current support for MITRE ATT&CK mapping or structured threat frameworks.

WSO2 – API Manager, Integrator, Identity Server, Asgardeo, Choreo, Bijira, Devant



Leader in



WSO2, founded in 2005 and headquartered in Austin, Texas, is a global provider of open-source and SaaS software for API management, integration, and identity and access management. Now part of the EQT portfolio, WSO2 has established a global presence with a customer base spanning 93 countries. The company's core mission is to simplify the development of digital experiences through a unified, AI-enhanced platform that caters to enterprises across diverse deployment models.

WSO2's strategic direction is centered around an integrated, extensible API management ecosystem that supports the full API lifecycle. Its vision builds on three foundational pillars: cloud-native, AI-augmented productivity; centralized governance across multi-gateway and multi-environment deployments; and open-source transparency with commercial-grade support. The company recently emphasized AI-native enhancements across its portfolio.

These include Bijira, a newly launched SaaS-native API platform, and an AI Gateway module designed to secure Generative AI interfaces.

WSO2's API portfolio includes several key components. The flagship WSO2 API Manager is a comprehensive open-source solution that supports the complete API lifecycle, including design, publishing, consumption, governance, and analytics. Bijira, built atop the Choreo cloud-native developer platform, delivers the same core functionality as a SaaS-native offering, enhanced with AI-powered features for productivity and governance. Other components include the Kubernetes Gateway, Universal Gateway, and Immutable Gateway, each optimized for specific deployment scenarios. Together, they form a federated gateway architecture managed through a single control plane.

API design is enhanced with AI-assisted tools that enable natural language input for generating API specifications, validating schemas, and producing mock implementations. This is supported by governance tooling that applies both static rule sets and AI-trained policies to enforce organizational and security standards. OpenAPI, GraphQL, gRPC, and asynchronous protocols are supported, with built-in specification validation and contract scanning for static analysis.

Discovery and inventory management are achieved through a centralized service catalog that spans all gateway types and environments. APIs can be categorized and tagged with business metadata, and the AI-powered Developer Portal assistant helps to quickly find appropriate APIs via conversational search. Support for multiple identity providers and organization-based access control further enhances visibility and contextual security.

API contracts are continuously validated against best practices, OWASP Top 10 threats, and custom organizational policies. Audit logs, revision tracking, and rollback capabilities contribute to a secure development pipeline. Runtime threat detection includes schema validation, request and response mediation, payload scanning, and protection against injections, schema poisoning, and other common web exploits. Partnerships with Salt Security, Ping Intelligence, and 42Crunch enable advanced integrations for anomaly detection, abuse prevention, and dynamic risk scoring.

WSO2 incorporates deep observability and analytics capabilities via both built-in dashboards and external integrations with platforms like ELK, Moesif, Prometheus, and Splunk. Behavioral anomaly detection, zero-day attack alerts, and token-based usage tracking are part of the standard analytics suite. AI-specific analytics further provide visibility into prompt and completion token consumption, cost management, and usage trends for LLM-backed APIs.

The platform's threat protection mechanisms include DoS/DDoS mitigation, schema validation, fine-grained access control with OAuth2, XACML, and OPA, and a library of rate-limiting and throttling policies. For GraphQL, additional protections such as query complexity and depth limits are provided. The AI Gateway extends these protections to LLM APIs, offering prompt templating, guardrails against data leakage, token-based quotas, and semantic caching.

WSO2 supports on-premises, hybrid, managed, and SaaS models with consistent functionality across environments. Its federated control plane allows customers to operate multiple gateways as managed nodes under unified governance. The platform is also well-suited to multi-tenant and B2B environments, with newly introduced support for organizational hierarchies and partner isolation models.

However, while the platform is highly extensible and functionally broad, it may require additional effort to configure deployments for advanced use cases compared to some competitors' turnkey offerings. WSO2 is working to bring parity between its software and SaaS versions, but currently the UI is inconsistent, and certain features, such as AI governance and guardrails, are available only in selected deployment modes.

WSO2 serves a broad customer base across government, utilities, finance, and telecommunications sectors. The platform is particularly well-suited for organizations seeking an open, extensible, and AI-enabled solution for API governance and security in complex multi-cloud or hybrid infrastructures.

Strengths

- Comprehensive open-source and SaaS offerings across API management, integration, and IAM.
- Unified control plane supporting multi-gateway federation and deployment flexibility.
- Deep API governance capabilities with both rule-based and AI-driven validation.
- Support for diverse protocols (REST, SOAP, GraphQL, gRPC, async) and API types.
- Centralized API analytics with real-time observability and custom dashboards.
- Integration with leading API security platforms.
- Fine-grained access control with role-based and organization-based scoping.
- Strong AI feature set including design-time assistance, chat-based discovery, and AI Gateway guardrails.
- Broad range of native DevOps integrations.

Challenges

- Full parity across software and SaaS offerings has not yet been achieved.
- Advanced features may require additional expertise to configure and deploy effectively.
- WSO2 has only recently added AI enhancements, such as semantic caching and model routing.
- Immutable gateway lacks support for certain AI features available in other gateways.

Vendors to Watch

Besides the vendors covered in detail in this document, we observe some other companies in the market that readers should be aware of. These vendors did not participate in the rating for various reasons but nevertheless offer a significant contribution to the market space.

Akana by Perforce

Akana is an API management platform originally developed by SOA Software and acquired by Perforce in 2021. Headquartered in Minneapolis, Minnesota, the company is known for its portfolio of developer productivity and DevOps tools. Akana continues to serve as the company's core offering for full-lifecycle API management, with a particular focus on regulated industries and large-scale enterprise deployments.

Why worth watching: Akana offers a comprehensive suite of capabilities for designing, securing, deploying, and monitoring APIs across hybrid and multi-cloud environments. Akana is well-positioned to support complex, long-lived API ecosystems where stability, scalability, and centralized control are critically important.

Apache APISIX

Apache APISIX is an open-source API gateway and microservices management platform developed by the Apache Software Foundation. The project has quickly gained global traction for its high-performance, extensible architecture and support for cloud-native environments. APISIX is written in Lua and built on top of NGINX and OpenResty, offering flexible plugin-based customization.

Why worth watching: APISIX is notable for its combination of performance, extensibility, and production-grade features that make it one of the most capable open-source alternatives to commercial API gateways. Native integration with Kubernetes and service mesh frameworks makes it especially attractive for multi-cloud and hybrid deployments. API7.ai, the company behind the project, also offers an enterprise solution based on APISIX with expanded capabilities.

Apiiro

Apiiro is headquartered in New York City, founded in 2019. Originally focused on software supply chain risk management, the company has since expanded into application security posture management (ASPM), offering a platform that provides visibility and governance across the entire application development lifecycle. With growing emphasis on API exposure as a critical part of software risk, the platform is now a key enabler of proactive API security.

Why worth watching: Apiiro brings a contextual, risk-based approach to API security by analyzing application source code, infrastructure configurations, and CI/CD workflows to uncover where APIs are exposed and how they can be exploited. Unlike traditional

perimeter-focused tools, it validates the security of APIs before they are deployed, identifying issues such as hardcoded secrets, missing authentication, and broken authorization flows during development.

Barracuda Networks

Barracuda Networks is a long-established security vendor based in Campbell, California, founded in 2003. Historically known for its email and network security appliances, the company has steadily expanded into cloud-native security offerings. Within its Web Application and API Protection (WAAP) portfolio, Barracuda now delivers capabilities that address growing API-related attack surfaces across hybrid and multi-cloud environments.

Why worth watching: The platform provides API discovery, schema validation, threat detection, and behavioral anomaly monitoring, available as a cloud-delivered service with centralized policy enforcement. Notably, Barracuda integrates protection into CI/CD pipelines, enabling shift-left security through OpenAPI specification analysis and pre-deployment validation. For organizations already leveraging Barracuda for perimeter or application-layer defense, its API security features offer a natural extension of existing infrastructure.

Boomi

Boomi is a cloud integration platform provider headquartered in Chesterbrook, Pennsylvania. Founded in 2000, the company was later acquired by Dell Technologies. Since 2021, Boomi again operates as an independent private company focusing on simplifying enterprise integration across applications, data, and systems. In 2024, Boomi expanded its API management capabilities through the acquisition of the Mashery platform from TIBCO. In 2025, the company further strengthened its position by acquiring APIIDA, a specialist in federated API governance, and relaunching its API management offering with a focus on agentic AI.

Why worth watching: With the combined assets of Mashery and APIIDA, Boomi has reemerged as a serious contender in modern API management, integrating federation, governance, and security across hybrid and multi-cloud environments. Boomi Enterprise Platform supports full API lifecycle management, productization, and analytics. Positioned at the intersection of integration, automation, and AI agent orchestration, Boomi's approach is explicitly engineered to support emerging AI use cases by exposing APIs as governed tools for autonomous agents.

Cerbos

Cerbos is an open-source software vendor focusing on adaptive authorization management. Founded in 2021, the company comprises a highly distributed engineering and management team, while being officially headquartered in London, UK. Cerbos has been focusing on developing an open source, language-agnostic dynamic authorization solution that is extremely lean and performant.

Why worth watching: Cerbos is an example of a vendor that only solves one specific problem but does it exceptionally well. The platform's high-performance, low-latency stateless implementation makes it especially suitable for securing access to business-critical APIs, but, of course, it can be used for a variety of other application development scenarios where dynamic authorization is required.

Citrix

Citrix Systems is a multinational software company that provides solutions for digital workspace, application delivery and security, and cloud services. Founded in 1989, the company is primarily known as a leading provider of remote desktop and application services, but as a part of its application security business, Citrix offers Application and API security – a comprehensive, layered security solution that combines web application firewall, bot management, API gateway, and SSL termination capabilities.

Why worth watching: available across multiple form factors (from hardware appliances to cloud-native containerized and SaaS offerings), the platform provides consistent, centrally controlled API protection across multi-cloud and hybrid environments and is available as a fully managed solution.

Data Theorem

Data Theorem is a company specializing in application security solutions. Founded in 2013 and based in Palo Alto, CA, the company offers a range of automated managed security services for developers of mobile applications and APIs. Through a large ecosystem of technology partners, Data Theorem offers a portfolio of SaaS agentless solutions for mobile, web, API, cloud, and supply chain security.

Why worth watching: Data Theorem offers an API security solution that spans discovery, risk assessment, and runtime protection, fully integrated into the company's broader Application Security Testing platform. Its solution continuously scans cloud-native environments to identify and inventory APIs, analyze their behavior, and detect potential misconfigurations or exposures. A notable feature is its automated generation of API-specific security policies based on real traffic analysis, reducing manual effort and accelerating time to protection.

Fastly

Fastly is an American edge cloud platform provider headquartered in San Francisco, CA. Established in 2011 as a content delivery network, the company has grown into a full-featured platform combining network services, compute capabilities, as well as security and observability functionality. In 2020, Fastly acquired Signal Sciences, a web application and API security vendor. Since 2022, it also owns Glitch, a popular online web application development platform.

Why worth watching: Fastly's security portfolio offers a broad range of capabilities including an intelligent next-generation web application firewall (WAF), API protection, DDoS mitigation, bot management, and managed security services. With its global edge cloud platform, Fastly can offer both scalability and flexibility for the most sophisticated deployments, as well as a high degree of automation for threat mitigation with its patented SmartParse technology.

HAProxy Technologies

HAProxy Technologies is a provider of high-performance load balancing and application delivery solutions, headquartered in Newton, MA and Antony, France. Known primarily for its open-source load balancer, the company offers HAProxy Enterprise as a commercial solution that extends core proxy capabilities with advanced traffic management, observability, and security. In recent years, it has been positioning its platform as a high-efficiency alternative for securing and routing API and AI traffic.

Why worth watching: HAProxy Enterprise offers a high-performance, modular API gateway with deep traffic inspection, fine-grained policy enforcement, and strong security controls. Its Global Profiling Engine and AI Gateway address both traditional API threats and emerging AI-centric use cases. With extensive customization and real-time updates, HAProxy offers a flexible, security-focused alternative for organizations requiring control, scalability, and integration with existing infrastructure.

IBM

IBM is a global technology and services provider headquartered in Armonk, NY, with a long-standing presence in enterprise integration, cloud infrastructure, and application security. Its API management solution, IBM API Connect, has matured into a platform supporting full-lifecycle API governance across hybrid and multicloud environments. With the recent acquisition of webMethods, IBM has further expanded its integration and API management portfolio, reinforcing its position in the enterprise connectivity landscape.

Why worth watching: IBM API Connect offers multiform API management with native support for REST, GraphQL, and event-driven APIs, enabling organizations to unify control over increasingly complex digital ecosystems. The platform addresses API sprawl with centralized governance, developer portal features, and monetization support. Its recent partnership with Noname Security introduces advanced runtime protection and threat analytics powered by machine learning. Combined with enterprise-grade deployment flexibility and integration with existing IBM technologies, API Connect remains a strategic choice for organizations requiring scalable, secure, and regulated API environments.

Link11

Link11 is a cybersecurity vendor headquartered in Frankfurt, Germany, known primarily for its expertise in DDoS mitigation and real-time application protection. Founded in 2005, the company has steadily expanded its service portfolio to address broader web and API

security challenges. In 2024, Link11 acquired Reblaze, an Israeli company specializing in cloud-native web application and API protection, marking a strategic shift toward a more comprehensive WAAP offering.

Why worth watching: Reblaze's strengths in API visibility, schema enforcement, and anomaly detection complement Link11's established reputation in infrastructure protection, allowing the company to offer a fully integrated, cloud-based platform for safeguarding both web applications and APIs. In addition, Link11's European presence ensures compliance with regional data sovereignty requirements.

Mayhem

Mayhem, formerly known as ForAllSecure, is a cybersecurity company headquartered in Pittsburgh, Pennsylvania and founded in 2012 with a focus on highly automated application security testing. Rebranded as Mayhem, the company has broadened its focus to cover modern development pipelines, with growing emphasis on API and service-level security testing.

Why worth watching: Mayhem applies intelligent fuzzing and autonomous testing techniques to uncover vulnerabilities in APIs and software binaries with minimal manual input. Its platform is designed to continuously analyze application behavior, identifying flaws such as memory corruption, logic errors, or unhandled exceptions in APIs. Unlike traditional security testing tools, Mayhem focuses on autonomous discovery of zero-days and subtle runtime flaws, making it particularly useful in high-assurance environments.

MuleSoft

MuleSoft is another veteran player in the API management market. Founded in 2006 in San Francisco, CA, MuleSoft has been focusing on providing a unified application integration platform to connect devices, applications, and data sources across on-premises and cloud environments.

Why worth watching: developing, publishing, and re-using APIs is the technological foundation for any integration platform, and the company provides a range of products and services for quick low-code development and testing of APIs, a comprehensive online marketplace for publishing and consuming APIs and other assets, as well as a data protection and security layer to stop threats and prevent data breaches.

Ping Identity

Ping Identity is a publicly traded software company headquartered in Denver, CO. Founded in 2002, the company has grown into one of the leading providers of identity federation and access management solutions. A leading provider of identity and access management solutions, the company has also expanded into API security by acquiring Elastic Beam, a pioneering security intelligence company.

Why worth watching: Ping API Intelligence is a real-time monitoring and threat detection solution for API traffic. By using automated API discovery and detection powered by AI models, the product can quickly centralize API monitoring, detect anomalies and suspicious activities, and block attacks automatically. Through integrations with other Ping products, it can offer comprehensive visibility and protection across on-premises and clouds.

PlainID

PlainID is a cybersecurity company headquartered in Tel Aviv, Israel, founded in 2014. Originally focused on policy-based access control for enterprise applications, the company has evolved into a leader in Authorization-as-a-Service. By abstracting authorization logic from applications and APIs, PlainID aims to simplify and centralize the management of fine-grained access policies across complex, distributed environments.

Why worth watching: As APIs become the backbone of modern digital ecosystems, controlling access to them is critical, and PlainID brings a robust, scalable approach to externalized authorization. Its Policy-Based Access Control engine enables dynamic, context-aware access decisions for APIs and microservices. The platform supports centralized policy lifecycle management, integration with API gateways, and enforcement via externalized decision points. This approach allows organizations to separate policy management from application logic, improving consistency, auditability, and speed of policy changes.

Radware

Established in 1996, with corporate HQ in North America and its international headquarters in Tel Aviv, Israel, Radware specializes in application delivery and cybersecurity solutions. The company provides a broad range of application security solutions in a suite, including API Protection.

Why worth watching: Radware API Protection maps the API attack surface by leveraging an automated discovery algorithm and generating tailored security policies to detect and block API-focused attacks in real time. It also uses a combination of access controls, DLP, bot management, and DoS mitigation tools to protect against API security threats.

Spherical Defence

Spherical Defence is a British security startup company based in London. Founded in 2017, the company is developing an innovative application security monitoring technology that is capable of unsupervised analysis of any machine-to-machine communications and JSON payloads - from HTTP traffic to system logs – analyzing over 150 telemetry points and detecting any anomalies in system or user behavior.

Why worth watching: As opposed to many other ML-based security solutions, Spherical Defense's product is fully autonomous and unsupervised – it does not require any manual configuration or training. It does not just identify anomalies in API traffic but can classify

them into multiple categories of attacks, including excessive data exposure, malicious injection, sensitive information transmission, and even adversarial attacks against ML models.

SmartBear

SmartBear is a software company headquartered in Somerville, MA, founded in 2003. Known for a broad portfolio of quality assurance and development tools such as Swagger, SoapUI, and TestComplete, SmartBear has long been a prominent player in the API lifecycle ecosystem. In 2023, the company expanded its API design and governance capabilities through the acquisition of Stoplight, a leader in collaborative API design and documentation solutions.

Why worth watching: SmartBear offers a uniquely comprehensive approach to API quality, covering the full lifecycle from design and mocking to testing, monitoring, and documentation. The platform supports multiple API specifications, including OpenAPI and AsyncAPI, making it suitable for modern microservices and event-driven architectures. For organizations looking to embed security and governance into every stage of the API lifecycle, SmartBear presents a compelling proposition.

StackHawk

StackHawk is a privately held cybersecurity company based in Denver, Colorado. Founded in 2019, the company has positioned itself at the intersection of application security and DevOps, with a strong emphasis on integrating security testing into modern CI/CD pipelines. StackHawk delivers dynamic application and API security testing as part of the software delivery process, making security actionable for developers.

Why worth watching: StackHawk is among the few vendors explicitly focused on developer-centric API security testing. Its core product is a modern take on dynamic application security testing (DAST), capable of scanning REST and GraphQL APIs for common vulnerabilities such as injection attacks, misconfigurations, and authorization flaws. The platform integrates seamlessly with build pipelines and source control systems, providing immediate, actionable feedback for developers.

Traceable by Harness

Traceable is a cybersecurity vendor based in San Francisco, California, founded in 2018. The company has quickly emerged as a prominent player in API security, offering a dedicated platform that combines deep observability with machine learning to detect and mitigate API-based threats. In early 2025, Traceable was acquired by Harness, a software delivery platform provider, aiming for a new standard in seamlessly developing, delivering, and securing applications.

Why worth watching: Traceable brings a data-driven, threat-focused approach to API protection, setting itself apart with its ability to automatically discover APIs, baseline their

behavior, and detect anomalies in real time. The platform supports runtime protection through behavioral analytics, sensitive data flow analysis, and fine-grained access enforcement. Its acquisition by Harness enhances the potential for tighter integration of API security into the software delivery pipeline, making it especially attractive to enterprises with large-scale, microservices-driven architectures in highly regulated industries.

Tyk

Tyk Technologies Ltd is a privately held company with sales offices located in London, Singapore, and Atlanta. Since 2015, it has been the primary force behind the Tyk Open Source API gateway and Tyk Enterprise, an API Management platform designed for DevOps. Comprising their own codebase built from the ground up instead of wrapping existing products from other vendors, the Tyk platform is designed for multi-DC and multi-cloud deployments, high performance, and full backward compatibility.

Why worth watching: Designed and maintained by a dedicated developer team, the open-source API gateway provides the full range of functionality free of charge, with commercial licensing available only for the management dashboard built on top of it. Tyk Enterprise includes an API management dashboard to manage, maintain and secure APIs across multiple gateways along with built-in policy management, operational analytics, and reporting. Tyk's integrated developer portal provides functions for developer onboarding, API documentation, and usage analytics.

Zuplo

Zuplo is a Seattle-based startup founded in 2021, focused on reinventing API management for modern development practices. Zuplo caters specifically to API developers by emphasizing programmability and developer experience. With support teams in North America and Europe, Zuplo serves a diverse customer base across various industries. Its flagship product offers a fully managed, cloud-native API management platform designed to be intuitive, extensible, and infrastructure-agnostic.

Why worth watching: By decoupling API delivery from centralized infrastructure and bringing APIs physically and logically closer to their consumers, Zuplo empowers organizations to transform APIs into monetizable digital assets. The platform's support for Git-based configuration management, instant deployment, and self-service onboarding lowers the barrier to entry for developers.

Related Research

[Leadership Compass: API Security & Management \(2023\)](#)

[Buyer's Compass: API Management and Security](#)

[Leadership Compass: Web Application Firewalls](#)

[Leadership Compass: Policy-Based Access Management](#)

[Advisory Note: The Role of APIs for Business](#)

[Whitepaper: The Dark Side of the API Economy](#)

[Executive View: Cequence Security Unified API Protection](#)

[Executive View: Curity Identity Server](#)

[Executive View: Forum Sentry API Security Gateway](#)

[Executive View: WSO2 Choreo](#)

[Executive View: Noname API Security Platform](#)

[Rising Star: Zuplo](#)

Copyright

© 2025 KuppingerCole Analysts AG. All rights reserved. Reproducing or distributing this publication in any form is prohibited without prior written permission. The conclusions, recommendations, and predictions in this document reflect KuppingerCole's initial views. As we gather more information and conduct more profound analysis, the positions presented here may undergo refinement or significant changes. KuppingerCole disclaims all warranties regarding the completeness, accuracy, and adequacy of this information. Although KuppingerCole research documents may discuss legal issues related to information security and technology, we do not provide legal services or advice, and our publications should not be used as such. KuppingerCole assumes no liability for errors or inadequacies in the information contained in this document. Any expressed opinion may change without notice. All product and company names are trademarks™ or registered® trademarks of their respective holders. Their use does not imply any affiliation with or endorsement by them.

KuppingerCole Analysts supports IT professionals with exceptional expertise to define IT strategies and make relevant decisions. As a leading analyst firm, KuppingerCole offers firsthand, vendor-neutral information. Our services enable you to make decisions crucial to your business with confidence and security.

Founded in 2004, KuppingerCole is a global, independent analyst organization headquartered in Europe. We specialize in providing vendor-neutral advice, expertise, thought leadership, and practical relevance in Cybersecurity, Digital Identity & IAM (Identity and Access Management), Cloud Risk and Security, and Artificial Intelligence, as well as technologies enabling Digital Transformation. We assist companies, corporate users, integrators, and software manufacturers to address both tactical and strategic challenges by making better decisions for their business success. Balancing immediate implementation with long-term viability is central to our philosophy.

For further information, please contact clients@kuppingercole.com.