White Paper

F5 Distributed Cloud API Security—Complete, Full API Lifecycle Security

Combine the power of data analytics and deep insights from AI and machine learning to discover, detect and protect your APIs. Easily identify and eliminate vulnerabilities, block API attacks, and prevent sensitive data leakage via API endpoints.

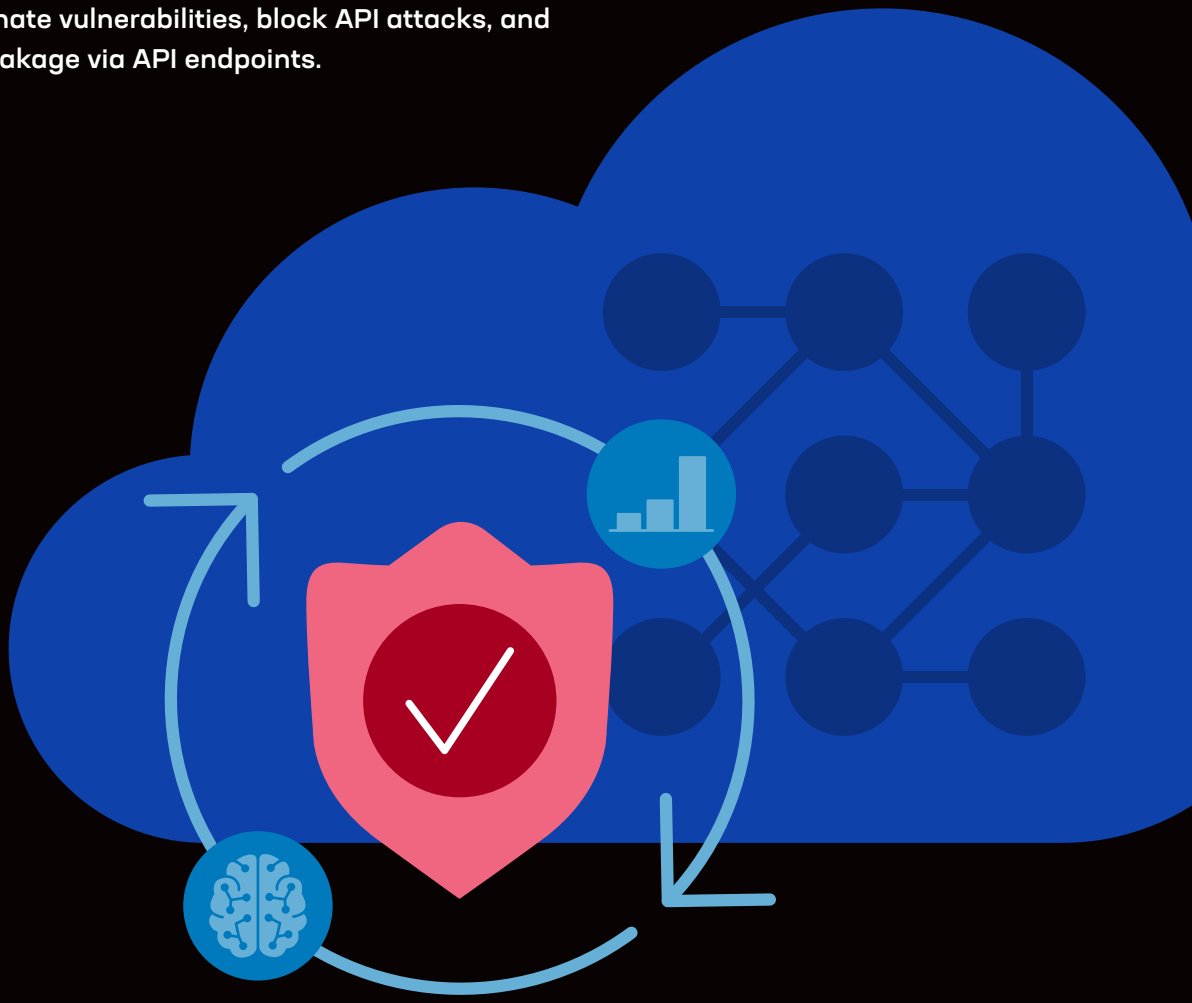

Key Benefits

Improve API Security
Combine complete, automatic API discovery and positive security with in-line enforcement capabilities, including WAF signatures, layer 7 DoS, rate limiting, IP reputation, allow/deny listing, and more.

Limit Data Loss
Better understand and monitor sensitive data being exposed by APIs and implement critical data exposure rules that limit and block from API responses.

Improve API Visibility
Observing API metrics from a single, centralized user interface for easy identification of all API endpoints and streamlined monitoring for anomalous or malicious activity, including zombie and shadow APIs.

Reduce Exposure of API Vulnerabilities
Begin discovery, testing and monitoring of APIs earlier in the API development lifecycle, improving visibility and understanding of vulnerabilities directly from code - limiting exposure in production.

Reduce Time Documenting APIs
Learn and generate OpenAPI spec (OAS) files to minimize manual tracking of all API endpoints.

Strengthen API Access and Authentication
Augment API gateway functionality, delivering enhanced visibility, oversight and control over API authentication and access. Identify gaps in API authentication, control access, and stop unauthorized attempts to exploit APIs and the back-end systems and data.

[1] F5 Labs, Post-Breach Analysis: Sophistication and Visibility, https://www.f5.com/labs/articles/threat-intelligence/post-breach-analysis-sophistication-and-visibility

Modern applications are challenging the traditional security paradigm. APIs represent a new and expanding technological lever for organizations' competitive advantages, driving speed to market for new digital capabilities. This represents a key growth driver. However, APIs represent an expanding attack surface, with new entry points to disrupt services and gain access to data, including Personal Identifiable Information (PII).

The pervasiveness of APIs and the unique role they can play in the security or vulnerability of any application, and thus an entire organization, can't be overstated. In an analysis of breaches[1] in recent years, F5 Labs noticed that in most API-related incidents, the breach method is technically very simple, impacting public-facing, poorly-secured API endpoints.

Security, when it comes to APIs, is easier said than done. With the wave of application security event data being generated for a growing number of applications, and with API endpoints being monitored by most organizations these days, it can feel like an impossible task to stay on top of everything. And as the pervasiveness of modern, microservices-based application development continues, so will the number of APIs. Thus, the application threat surface will continue to get more difficult for organizations to deal with.

Why Are APIs Vulnerable?

There are many reasons why API security is difficult. First, consider the sheer scope and complexity of modern application environments and APIs deployed across highly distributed multi-cloud and hybrid architectures. This is all compounded by the speed at which apps and APIs are being developed. New connections and services are being introduced, including updates to existing APIs, through rapid CI/CD development cycles.

APIs are developed with common transport protocols REST, GraphQL, and more, which can contain flaws, creating vulnerabilities that can be exploited just like the applications they serve. It's likely most enterprises don't know all the APIs running in or connected to their environments. API visibility is a blind spot for many organizations. When speed and innovation are the goals, rigor in API documentation and tracking is often not the focus for developers.

Moreover, organizations with acquisition and integration activity, where IT environments and applications have been inherited, are dealing with situations where there are unknown applications, and APIs that are not well documented. This can create a large security blind spot. APIs are fast, lightweight, and reliable, often enabling critical communications and transfer of data between applications or clients. They have the potential to expose sensitive data, so they have become a desired target. This ever-growing, complex threat surface provides a struggle for legacy security technology and operations teams, pushing more of them to the brink of what is possible to try to secure and keep up with.

What Core API Security Capabilities Do Organizations Need to Implement?

In F5's latest State of Application Strategy Report for 2024, organization with revenue between $200 million and $1 billion reported having an average of 235 apps but almost 500 APIs.

APIs are quickly expanding the threat surface for most organizations. F5's latest research shows that on average, deployed APIs significantly out number apps. And this number tends to increase with company size, so if an organization expects to grow they should expect API proliferation, for example organizations with revenue between $200 million and $1 billion reported having an average of 235 apps but almost 500 (499) APIs.[1] Historically organizations have deployed API gateways as a valuable layer of protection, handling elements of API security. This included versioning and publishing, schema validation, monitoring, connectivity and routing, access, authentication, and rate-limiting. In F5's latest State of Application Strategy Report for 2024, 95% of respondents reported that their organization uses an API gateway. However, this is simply the first step and bare minimum when it comes to API security. Organizations need to treat API security more comprehensively, just as they do their core web applications. There are core elements of API security that organizations should be prioritizing when it comes to their application security stack. This includes technology and services that can provide continuous **API visibility and discovery from early on in the development life-cycle and throughout production**. Relying on securely developed and well-documented APIs with schema enforcement functionality (Positive Security) is critical, but only part of the equation. Organizations need capabilities to constantly learn, map and test their APIs including those that may have been forgotten, or they do not own and aren't documented, across all communication paths of an application. API discovery and testing maps API landscapes from source code and production, revealing inventory gaps, shadow APIs, abandoned 'zombie' APIs to block, and legitimate APIs needing governance—ensuring accurate inventory and oversight for effective security.

Knowing an API exists and having access control capabilities are two critical pieces to the API security puzzle. In the F5® 2022 State of Application Strategy Report, 68% of respondents ranked authentication and authorization as the most valuable components of API security. Not far behind was behavioral **analysis and anomaly detection** to monitor APIs in production (runtime), identifying and alerting on abnormal behavior and potential abuse, since there are ways to bypass authentication and authorization. Being able to track API behavior over time should include API request analysis and time series anomaly detection to build baseline behavioral attributes that can be used to identify anomalies in API request rates, errors, latency, throughput, and more. With this functionality, an alerting element is critical to raise issues when unexpected spikes or drops occur, unique traffic patterns are present, or abnormal API requests are detected.

Rounding out a modern API security stack requires an **in-line application and API security enforcement engine**. This most likely includes a WAF with multiple layers of application security functionality, such as granular application layer 7 policy enforcement with rate limiting, IP reputation, allow/deny list functionality, and layer 7 DoS. All this gives organizations the capabilities to further investigate and act on malicious endpoints, users, and other activity. This allows operations teams to quickly and easily identify suspected API abuse, plus act as anomalies are detected.

F5® Distributed Cloud API Security delivers a comprehensive approach to protecting API's across their full-lifecycle from build and test through to release, operate and monitor. The service allows organizations to more effectively **discover** all their APIs including manage their inventory and documentation, to more consistently **detect** behavioral anomalies and vulnerabilities within APIs, and **protect** their APIs with continuous inspection, unified policies and schema enforcement. Let's explore how.

Attacks Attempting to Leak or Exfiltrate Data via APIs

APIs are often implemented in a default or generic way without considering implications to the data they are intended to collect or connect to and the individual sensitivity of different data types to an organization, especially when certain compliance requirements are relevant. Often data is sent or exposed within APIs unknowingly or inadvertently. That's why being able to identify web app and API endpoints, where potential personal identifiable information (PII) and other sensitive data is being transferred or exposed, is critical, so it can be protected, and breaches prevented. Distributed Cloud API Security helps organizations get a handle on their API landscape and provide a view into sensitive data being exposed via their web apps and APIs, helping them better secure business and customer data, plus document and better understand their compliance posture.

Organizations can easily configure sensitive data policies to discover, tag and report on critical data being exposed within their APIs. This includes basic policies to identify common PII data (e.g. credit card numbers, physical and email addresses, and phone numbers), to specific compliance frameworks that can be applied with hundreds of predefined data types relevant to over 20 critical compliance frameworks (e.g. PCI-DSS, HIPAA, GDPR, SOC2 etc.) and even custom sensitive data unique to an organization. Once configured the service will automatically discover and document an organizations APIs directly from code repositories, analysis of traffic (requests and responses) and through client-side web crawling with the ability to view all endpoint details for each individual API. This includes the detection and flagging of PII that is being exposed and tagging of the relevant compliance framework(s).

Distributed Cloud API Security also includes a custom sensitive data detector (regular expression/regex based) functionality, allowing users to specify and search for less common or unique patterns which may be indicators of other sensitive data types in API requests and responses. This can be used to search for any unique, organizational specific data that needs to be detected and protected, including active monitoring of all API traffic to spot any inadvertent leaks or suspicious activities. These critical detection capabilities allow organizations to quickly and easily identify any critical data that is being shared for any API, so that remediation can be put in place.

Limiting and masking sensitive data over APIs is crucial for protecting data privacy, complying with regulations, reducing the attack surface, preventing unauthorized access or disclosure, and mitigating the impact of potential vulnerabilities. It is an essential practice to ensure the security and trustworthiness of API implementations. The Distributed Cloud platform has a variety of capabilities to help organizations protect sensitive data that is identified within web apps and APIs including Dataguard functionality that masks data in HTTP/HTTPS responses using a string of asterisks (*). Specifically for APIs the service includes Secure API Guard functionality, which delivers sensitive data masking and leakage detection capabilities for APIs. This allows organizations to establish API data protection policies, defining how data is handled within API responses to limit, block and/or mask. These policies controlling exposure and masking of data within APIs can easily be applied to specific API endpoints, a group of endpoints, specific paths or an entire domain (e.g. all APIs), ensuring that even if an attacker gains access to a given APIs traffic, the sensitive data remains secure and incomprehensible—rendering the data useless. On top of the masking capabilities, the service also includes continuous monitoring of all APIs with analysis of all transmitted data to help detect and report on any inadvertent leaks or suspicious activity of data within API responses.

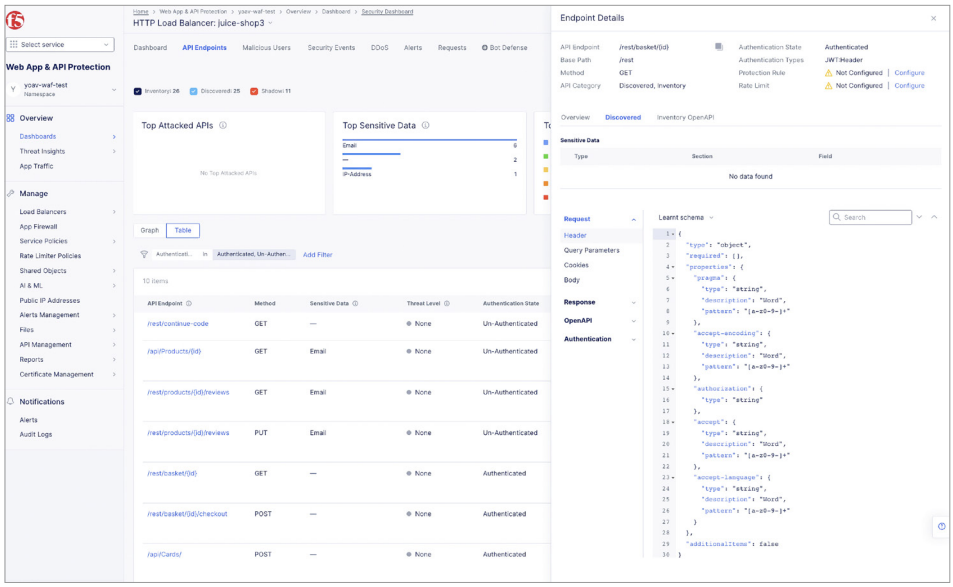

Figure 1: Learned API schema with ability to drill down into usage baselines and view any PII information at the individual API level.

Resource/DoS Attacks and Abuse of APIs

Like with any network or compute resource, APIs are susceptible to abuse and denial of service (DoS) attacks. APIs respond to client requests with responses which require CPU, memory, RAM, and more, with resource consumption being dependent on logical processing of inbound requests or the amount of data returned. Without rate limiting or other layer 7 DoS protections in place, this leaves web apps and APIs vulnerable to a single user or group flooding an API endpoint with too many concurrent requests. Such activity can slow down the service, leave an API unresponsive and, in many cases, lead to a denial of service of an endpoint.

It's critical that organizations implement or deploy technology that can provide rate limiting and other DoS mitigation functionality at layer 7 for web apps and APIs. The Distributed Cloud platform has the layer 7 DoS capabilities and rate limiting functionality to ensure service availability of web apps and APIs.

Organizations can granularly control API endpoint connectivity and the rate of requests. They can identify, monitor, and block specific clients and connections all together or set particular quotas or thresholds (number of requests allowed) with a duration (over a set period of time), and discrete HTTP methods to be rate limited. This granular control of API connections and requests can be done for individual APIs or an entire domain.

On top of this granular rate limiting functionality, the Distributed Cloud platform delivers robust layer 7 DoS attack detection and mitigation for web applications and APIs using a combination of techniques that includes alerts and blocking from traditional signature-based WAF functionality as well as anomaly detection and alerts from AI/ML. Machine learning happens in the centralized control plane, using metrics and log data collected from an organization's endpoints.

The visibility generated by this continuous analysis and baselining of web app and API behavior provides practitioners and organizations with the insights and alerting on anomalies, which can be used to generate layer 7 protection policies. There are a variety of remediation actions that can easily be put into place. These include rate limiting or deny listing based on IP address, region/country, ASN or TLS fingerprint, plus more advanced rules defining specific match criteria guiding app and API interactions with clients, including HTTP method, path, query parameters, headers, cookies, and more.

Combined, the rate limiting and layer 7 DoS functionality, along with behavioral AI/ML capabilities, of the Distributed Cloud platform delivers a rich set of capabilities to keep web app and API endpoints free from abuse, running smoothly, and available to process requests.

Injection Attempts and Other App and API Vulnerabilities

Monitoring API traffic in production is important, but it provides a reactive view of an organizations security, capturing issues only as they appear. For organizations, this means unknown vulnerabilities could persist, leaving critical services, business processes and sensitive data at risk. Solely relying on traffic analysis alone often means risks are identified too late, resulting in persistent exposure, potential breaches and service disruptions that impact organizational trust and operational stability.

Within Distributed Cloud API Security, we've integrated API testing into our multi-lens discovery capabilities, enabling proactive security testing across your entire API inventory. This automated approach identifies hidden vulnerabilities throughout the API lifecycle, including endpoints in pre-production environments. Our targeted security tests cover the complete OWASP API Top 10 threat categories, detecting unauthorized access, authentication flaws, sensitive data exposure (including PII), and misconfigurations. Organizations gain comprehensive visibility into their API threat landscape with actionable insights to protect endpoints, services, and data before vulnerabilities can be exploited. By implementing continuous, pre-emptive testing, you can ensure critical API services remain secure and stable without relying solely on production traffic to reveal threats—when it may already be too late.

Once in production, traditional WAFs still play a huge role in the protection of modern applications and the APIs that drive them. APIs are susceptible to the same types of injection attacks as the applications they support, including injection flaws like SQL, NoSQL, Command Injections, and more, attempting to execute unintended commands or access data.

The Distributed Cloud platform includes F5's core WAF with it's robust attack-signature engine containing over 8,500 signatures for CVEs, plus known vulnerabilities and techniques identified by F5 Labs. The service also includes threat campaign functionality which delivers protection against sophisticated, multi-vector attack campaigns by using fully vetted attack campaign signatures developed by F5 threat researchers. These signatures help protect organizations from a variety of injection attacks and other critical vulnerabilities and attack types, including DoS, bots, and automation, which are trying to exploit vulnerabilities that exist in the underlying code.

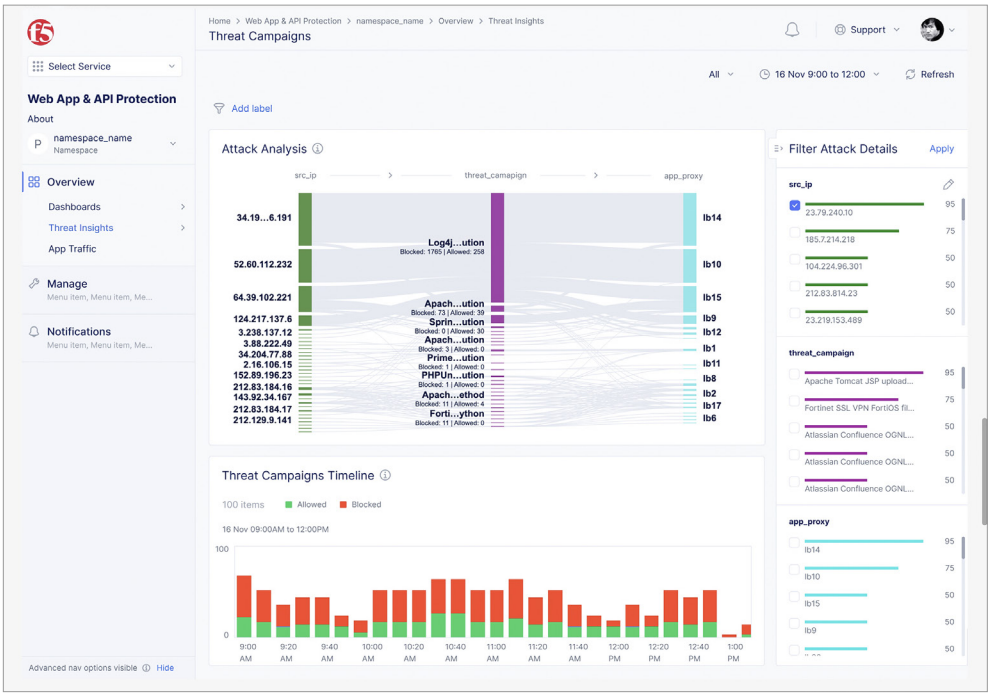

Figure 2: Threat Campaigns correlates singular attack incidents as extensive and sophisticated attack campaigns developing signatures to protect web apps and APIs from persistent attempts to exploit their code.

The WAF engine and Threat Campaigns are yet more elements that the Distributed Cloud platform delivers to protect applications and their APIs from exploitation, tightening up critical vulnerabilities in code while development teams work to review, patch, and improve code over time.

Gaps in Access and Authorization of APIs

When it comes to handling access and authorization threats, Distributed Cloud API Security has several capabilities that augment API gateway functionality, delivering API testing, enhanced visibility into API vulnerabilities, oversight, and control over API behavior, authentication, and access. This helps organizations identify gaps in API authentication, control access, and stop unauthorized access attempts to APIs and the back-end systems and data they connect. The service learns, models, and maps all app and API endpoints, including the status and type of authentication present.

API Authentication Discovery helps organizations automatically identify and baseline the authentication state of all APIs within an environment. The service can learn and document authentication types along with other API endpoint details based on direct code analysis, and traffic based discovery. Allowing organizations to better understand and easily associate authentication information with individual API endpoints for examination, and in support of a positive security model to maintain appropriate authentication of their APIs and in the development of critical API protection rules to limit or control API behavior and access. With OpenAPI spec files, either learned or uploaded the authentication information, which is part of OpenAPI spec details, can be automatically enforced, and unauthenticated traffic can be stopped at the edge, removing the need for origin API gateways and servers to handle these requests.

The service also includes JWT Validation functionality which allows organizations to upload authentication keys and will validate JSON Web Token (JWT) sign in requests at the edge. With this capability organizations don't have to store session states on the server and load user information from a database or cache. This immediate validation negates the need to go back to the origin for verification and increases the scalability of APIs—providing an overall faster client-server experience.

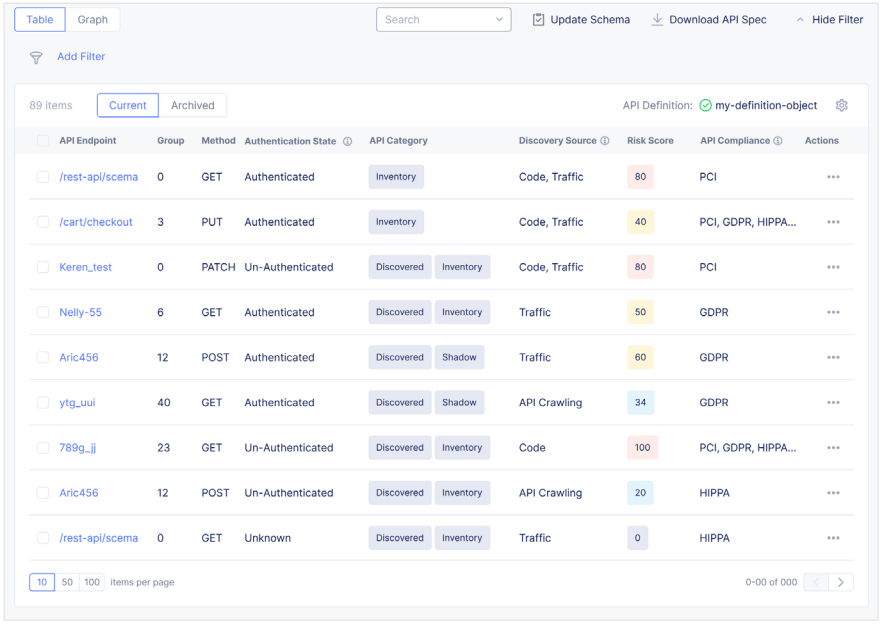

Figure 3: API Authentication Discovery and Validation–discover and view authentication status, plus other API details including sensitive data and compliance status, API category and risk score for all APIs including the ability to create protection rules.

This also contributes to improved security by allowing edge servers to filter out unauthorized requests before they can reach an organization's origin infrastructure. RSA private/public key pairs for JWT signature verification are supported, ensuring that the data in JWT payloads has not been modified by third parties. An identity provider first signs the JWTs by using a private key and Distributed Cloud API Security then verifies the integrity of the JWT by using the public key that is uploaded.

For endpoints with JWT authentication type, Distributed Cloud API Security can evaluate existing authorizations within tokens. This further helps organizations understand the security of their API posture, enabling swift action to take place if and where gaps are identified. The service can discover and validate headers, payloads, and signatures within JWTs which may be indicators of compromise. It can detect user role or user ID, identifying sensitive data in JWT payloads, ultimately producing an API endpoint threat level and risk score, which can be used to guide remediation efforts to shore up insecure endpoints. The API threat level developed for JWTs is composed of a variety of inputs, including these potential vulnerabilities:

- Digital signature algorithm is missing
- Expected signature algorithm is not enforced
- Inadequate JWT expiration policy
- Expired tokens are accepted
- Tokens with invalid signature are accepted
- Sensitive data found in JWT token

This discovery and analysis of authentication within API endpoints delivers remediation insights to organizations, so they can easily view the distribution of authentication types across API endpoints, including those endpoints where authentication isn't in place. The insights are delivered with rich filter and drill-down capabilities, so that operations and developer teams can act quickly react via existing API gateway(s) by updating authentication and access policies or through implementation of layer 7 policies to block or limit unauthenticated clients with Distributed Cloud API Security.

API Access Control

Another key element is the control of access to APIs, and this can be accomplished by implementing a positive security model. This is possible with Distributed Cloud API security, leveraging imported or discovered and automatically generated OpenAPI spec files and API protection rules (layer 7 policies). These are used to define specific API endpoints, API groups or Base paths, then set policies to enforce granular API access control, enabling a positive security model for applications and APIs. Any learned, generated OpenAPI Spec files for discovered APIs, API definitions and groups can also be exported into an existing API gateway where authentication and enforcement can be applied and appropriate API behavior can be monitored. These capabilities help organizations ensure only API connections and client requests to operations that are specified in their OpenAPI spec files are allowed and all other requests, including those where users are trying to guess objects, parameters or other unspecified operations, are denied.

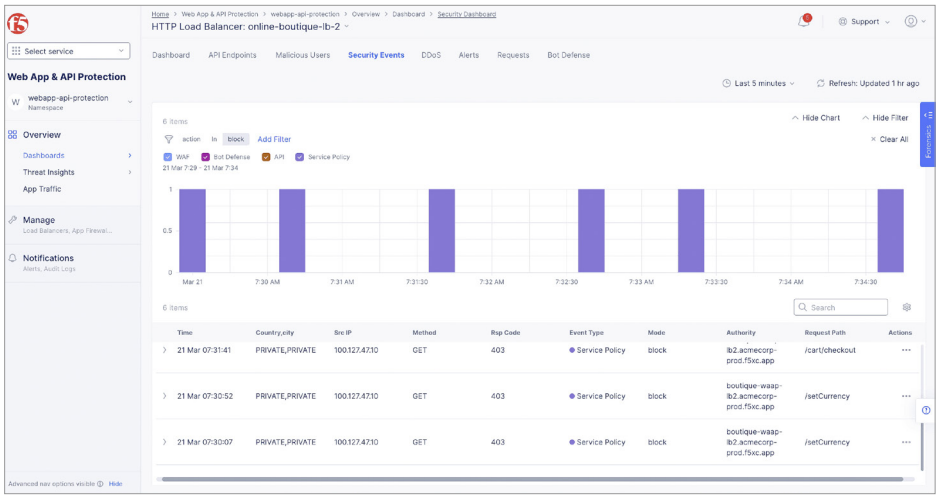

Figure 4: Protected API endpoint where request is returning 403 response code to client(s) who are not authorized to access.

If there are applications with flaws in the design of their access and authentication for any of their APIs, Distributed Cloud API Security with layer 7 protection policies can be configured to allow or block clients and their requests, thereby preventing attempts to bypass access and authentication functions.

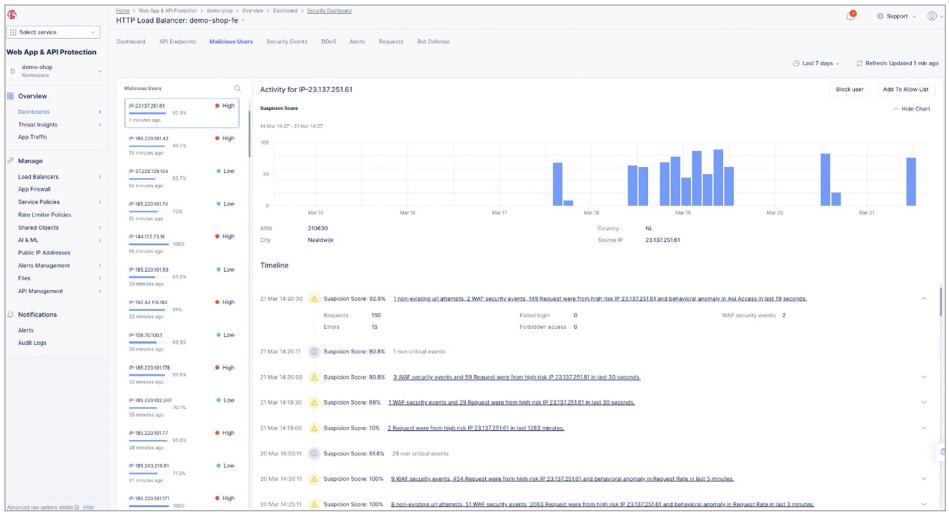

Figure 5: Malicious User functionality–correlates WAF events, abnormal API access behavior, failed login attempts, and so forth, for app and API endpoints, flagging these for operations teams with the ability to quickly take action including block/allow list, rate limit, and more.

Controlling Access with Behavioral Based Detection and Mitigation

The Distributed Cloud platform has a variety of AI/ML powered capabilities that monitor and help organizations enforce proper access and authorization of applications and APIs. One of the core features is malicious user detection and mitigation, which is backed by the platform's core ML engine. It monitors all client interactions with an application, including those of APIs, and analyzes them over time to identify behavioral outliers. Based on this analysis, each client is given a risk score (high, medium, or low) relative to all their interactions with a given app and any API endpoints.

Based on a client's activities with respect to a set of problem categories, a client's threat level will rise or fall. The ML engine learns client behavior from traffic generated across all endpoints in an organization's environment. There are various methods for malicious user detection, including forbidden activity via configured layer 7 protection policies (including HTTP methods, paths, query parameters, headers, and more), failed login attempts, WAF and threat campaigns signatures, bot defense and automated events, IP reputation database triggers, rate-limiting events, and invalid or nonexistent URL/request activity (resulting in 404 not found response codes).

When it comes to API security and protecting against unauthorized access via APIs—either through credential stuffing, brute force, or other forceful login attempt mechanisms—the ML engine can help by identifying failed login attempt activity or attempts to discover API parameters, and flag that to operations teams. The feature keeps track of the number of login attempts that have failed (specifically 401 unauthorized response codes) for all clients. When the number of login failures from a client exceeds the limit set, and/or there is a large spike based on historical, learned behavior, the client will be classified as malicious and can be blocked. The behavioral-based app and API monitoring and mitigation capability of malicious user detection is augmented with some other features specific to detecting anomalies in API access and session flow.

API Access Anomaly Detection uncovers endpoint abuse by looking for irregular traffic across multiple endpoints; understanding each API request in the context of the user making it, their previous requests, and the overall request patterns of all users across all APIs; and leveraging behavioral analytics and machine learning (ML) to detect abnormal usage patterns in APIs.

Distributed Cloud API Security is bolstered with these ML-powered capabilities which can help organizations protect their APIs including validating connections and access, monitoring behavior and alerting on anomalies over time, and helping to identify unusual client behavior to pinpoint potential areas of compromise. Not only will the service test and report on API authentication mechanisms and potential access issues, but it includes critical enforcement mechanisms to easily act (e.g., set a block/allow/limit choice) against unwanted client or API activity quickly. The service can play a significant role in modern application security by augmenting the efforts of API gateways in controlling client interactions and enforcing appropriate access to API endpoints.

Improper Asset Management and Security Misconfiguration

Another critical component of API security is ongoing management of APIs. This can take many forms. Some of this is handled by API gateways, but modern web app and API protection capabilities are necessary to aid in this, as modern applications driven by APIs can be very dynamic with quick and rapidly evolving development cycles. It can be hard for operations, security, or IT teams to keep up, as modern apps with APIs tend to expose more endpoints than traditional web applications and these APIs are usually changing frequently with new ones being added at a rapid pace for many organization today. This makes having well-documented, on-going visibility and hygiene of APIs important. Properly documented and deployed API versions, inventory management, and tracking, plus robust, continuous API discovery and testing, play an important role in protecting an organization and mitigating issues such as old, deprecated APIs or API versions, and unknown or shadow APIs and OWASP API Top 10 vulnerabilities. Distributed Cloud API Security allows organizations to easily deploy and enforce a positive security model when it comes to protecting their APIs. This ability, combined with complete web app and API path discovery, testing and automatic generation of API documentation, helps identify, model, and baseline unknown or undocumented APIs, plus identify critical vulnerabilities that require their attention.

Dynamic API Endpoint Discovery and Schema Learning

Not all APIs are known or well documented and within many organizations there may be multiple teams developing services and APIs across a variety of different environments. As modern app environments grow more complex, and the number of APIs grows, it's imperative for organizations to keep up, ensuring they have a consistent and complete understanding of their API endpoints and any vulnerabilities. It's imperative to get this visibility as early in the software development lifecycle as possibly, before apps and APIs are released into production, limiting exposure time of potential vulnerabilities. As a result, organizations should look to maintain centralized, continuous discovery, consistent documentation and monitoring of all web app and API endpoints that are present in their environments.

To help organizations stay on top of their APIs, as software development cycles continue to evolve, and speed up – complete API discovery and continuous API testing is paramount. This includes code-based discovery, combined with dynamic API discovery from traffic and intelligent web crawling that will continuous crawl web-based apps from the client-side - plus automatic OpenAPI spec file generation, which are all part of Distributed Cloud API security. The service supports discovery for a variety of protocols including REST, SOAP, gRPC, and GraphQL. Organizations can easily enable code-base discovery across a variety of developer repositories (e.g. Azure, BitBucket, GitHub, Gitlab and more) to scan, map and document their APIs and any vulnerabilities earlier on in the development lifecycle. There are a variety of supported languages which can be scanned including Java EE, Java Spring, .NET, Python

flask, Python Django, Javascript express, JavaScript hapi and GO. This initial inventory, mapping and documentation done can be paired with continuous traffic-based discovery.

By analyzing production traffic, looking at API requests and responses the service provides another lens and more complete picture of an organization's entire API threat surface. In addition to traffic-based analysis for API discovery, the service also provides a client-side crawling capabilities for API discovery. This is a discovery technique that independently simulates user behavior to detect exposed APIs within web applications. It focuses on identifying other APIs that are not discovery via traffic analysis or within scanned code basis, providing a more comprehensive view of your API landscape. External crawling is great to uncover old APIs that are no longer in use, but left exposed, or infrequently used APIs which may not be consistently passing traffic, thus is complimentary to traffic-based discovery. These options for traffic and crawler-based API discovery can be enabled standalone as the first step to map and document an organizations APIs or used as an additional lenses of discovery to augment what was learned through analysis of the code repositories or already documented and in inventory via imported OpenAPIspec files to uncover unknown, shadow APIs, unused/Zombie APIs and any other discrepancies including old or deprecated APIs, versioning issues and intended versus actual API behavior. With this discovery functionality comes critical testing and inventory management capabilities, allowing users to continuously test their APIs in inventory and identify vulnerabilities, plus maintain a clearer picture of their API threat landscape. For defenders, managing their API inventory is crucial and this service allows them to seamlessly group and tag APIs plus move endpoints to and from inventory, helping them maintain a clearer picture of their APIs. Users can effortlessly promote discovered or shadow API endpoints into a centralized inventory, while false-positive discoveries can quickly be marked as "non-API" when they are mistakenly detected and removed from inventory. This allows organizations to maintain a clean and precise inventory with insights into API vulnerabilities, and documentation reducing unnecessary clutter and making it easier to manage security and enhance the accuracy of protections applied to all API endpoints.

Positive Security with Import and OpenAPI Spec Validation

Organizations can leverage Distributed Cloud API Security to enforce proper API behavior based on valid API definitions through automatically generated or imported OpenAPI specs. Documented API characteristics are used to validate input and output data from API endpoints like data type, minimum or maximum length, permitted characters, or valid values ranges. The service will check API traffic for compliance, allowing organizations to automatically validate API traffic and block or implement protection rules, further limiting or controlling access to individual API endpoints, API groups, or base paths defined in the spec file. This includes the enforcement of undefined parameters, allowing users to specify whether to block or allow requests containing parameters not explicitly defined in an OpenAPI spec file.

Behavioral Analysis Through Machine Learning

Once in production, the Distributed Cloud platform has the capabilities to discover and perform behavioral analysis on the various logs collected from endpoints and APIs of an application using ML. The schema structure of APIs and authentication elements is learned by analyzing sampled request and response data examples for each API. This provides learning of the behavior of these paths, including request and response schemas and sensitive data detection. The discovery and mapping capability shows the complete inventory of learned application and API paths, including shadow sets, with development of an OpenAPI spec for exporting and usage by development teams. Downloaded OpenAPI spec files for learned API schemas can be rendered at a variety of levels, including HTTP load balancer, app type and per API. The file is an exportable JSON file.

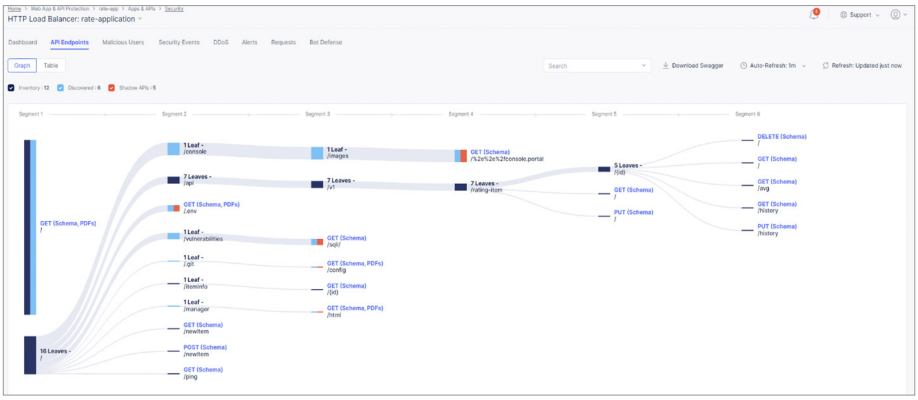

Figure 6: Learned app and API path visualization includes Inventoried APIs, Discovered APIs and Shadow APIs. The endpoint paths are shown in a hierarchical structure with root and leaf relationships presented in segments.

As application paths and API endpoints are discovered and monitored, probability distribution functions related to each endpoint are generated for metrics such as request size and response size, latency with and without data, request rate and error rate, and response throughput. Analysis is performed periodically, and these baseline metrics are updated. Learning of the API endpoints and associated metrics is incremental in nature and updated periodically. This visualization and the corresponding metrics can be used to create layer 7 security policies to control access to and functions of APIs. These resulting rules can be applied to all APIs, including unknown/shadow APIs, to block or limit activity.

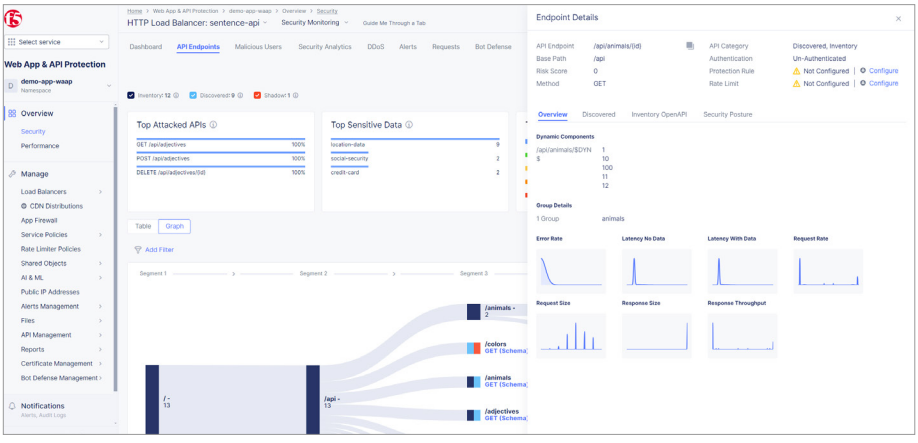

Figure 7: Endpoint details include metrics such as request size and response size, latency with and without data, request rate and error rate, and response throughput.

There are a variety of rule types that can be applied to web application paths and APIs when using the Distributed Cloud platform. They include:

- **Protection rules for APIs (allow/deny, rate limit, and more)—**allow/deny or rate limit API endpoints, using API endpoint path and method(s) for applying protection rules
- **Update/configure API request parameters—**HTTP query parameters, HTTP headers, web cookie and/or TLS fingerprint match criteria/parameters to limit functionality of specific API endpoints
- **IP reputation—**client access can be allowed/denied based on IP reputation categories—all F5 categories, specific groups, or individual IPs—so organizations can select options to match by IPv4 prefix, IP prefix, ASN list, or BGP ASN sets

Not only can organizations easily discover unknown or shadow APIs, as well as upload and enforce positive application and API security, but Distributed Cloud API Security can deliver more holistic visibility with a comprehensive API endpoints dashboard. This includes all discovered and imported API endpoints for each domain, with views into Top Attacked APIs, Top Sensitive Data, Total API Calls, Most Active APIs, plus risk and threat scores for each API endpoint in a given domain. This centralized view allows operators or engineers to move quickly, identifying potential issues within their API environment, with the capability to easily drill down, investigate, and take action as appropriate to neutralize any anomalies or threats that could impact connectivity, availability, or app and API security.

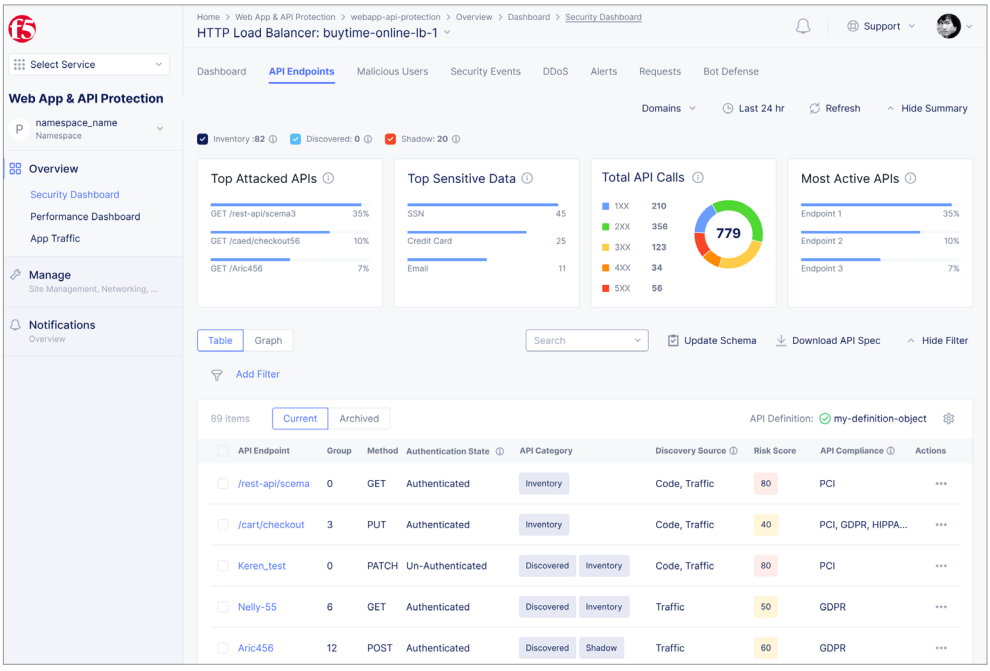

Figure 8: API endpoints dashboard—all discovered and imported API endpoints for each domain with views into Top Attacked APIs, Top Sensitive Data, Total API Calls, Most Active APIs plus risk and threat scores for each API endpoint in a given domain.

Conclusion

Deliver superior digital experiences with performant, effective, and scalable application and API security with F5

Applications and increasingly APIs are the lifeblood of most businesses today. Modern applications call for scalable and adaptive security solutions. As applications become increasingly modular, complex, and distributed, they require security services that can do more. Distributed Cloud API Security as part of F5's Distributed Cloud WAAP delivers the cybersecurity efficacy and ease of use that today's application architectures require. It's a better way to secure modern applications and APIs with unparalleled performance and availability at scale, offering consistent operations, security, and end-to-end observability.

With Distributed Cloud WAAP, organizations benefit from comprehensive application and API security to more effectively secure and manage APIs throughout their full life-cycle, helping drive business velocity by enabling extensive, modern application and API deployments with the necessary oversight and protection against API-specific threats. Organizations can seamlessly augment existing API management and gateway functionality to identify and secure API vulnerabilities, access, and authentication gaps, enabling API security for all critical threat categories. With Distributed Cloud API Security, customers can pair complete API discovery and mapping to identify unknown or shadow APIs, with the necessary testing, detection and enforcement tools, in one global, SaaS-delivered solution.

This innovative and accessible service helps reduce app and API security gaps and enables consistent coverage across an organization's entire application portfolio. With Distributed Cloud API Security, organizations can simplify their path to effective security while fostering the innovation their business and customers demand.

Explore more and request a free trial at f5.com/cloud/products/api-security.

[1] 2024 F5 State of Application Strategy Report