



Choose the right deployment architecture when securing your apps with Distributed Cloud Web App and API Protection (WAAP), combining security with the flexibility to deliver apps where and how they are needed. Extend application and API security across hybrid, multi-cloud, and edge environments.



KEY BENEFITS

Get SaaS security for all environments

Protect web apps and APIs in any hybrid, multi-cloud, edge, or on-prem environment with a single SaaS security solution.

Deploy anywhere with a purpose-built solution

The F5® Distributed Cloud Platform is built from the ground up to deliver advanced, accessible security for modern, highly modular, distributed applications.

Empower SecOps to do more, faster with end-to-end observability and policy enforcement

Improve your SecOps team's efficiency by providing visibility and unified security policies that are portable across clouds and premises, with end-to-end observability and policy enforcement.

Lower TCO with integrated services, SaaS, and flexible deployments

Reduce overall Total Cost of Ownership (TCO) by collating disparate cloud security solutions onto a single cloud security stack that eliminates the need to maintain separate policies and infrastructures.

Enhance developer experience and speed up time-to-service

Optimize the developer experience by easily plugging into existing CI/CD workflows and DevOps tools, allowing automated security workload deployments and validation to be part of your normal processes.

Extend Application and API Security Across Multi-Cloud and Edge Environments

Application workloads are increasingly deployed across multiple diverse environments and application architectures. Modern apps have some distinct characteristics that make their management more difficult or different than traditional (monolithic) apps. As a result, their security becomes more complicated.

- They are often designed to be modular, built with APIs
- They are built with agile build methodologies, so are constantly evolving/changing
- They are often much more segmented and complex
- They can be highly distributed across different environments

These attributes contribute to increases in performance and scale, improving speed to delivering innovation, greater flexibility, and streamlined maintenance. However, they also introduce new challenges in the overall visibility of app performance and security, specifically in the implementation of application security. There are constant trade-offs being discussed and made within organizations about how security should and is being handled, based on their ever-changing apps, customer experience, and the complexity of application environments as they evolve.

Organizations don't want security to be a hinderance or cause unnecessary friction of app evolution, customer experience or performance. When organizations think of application security deployments in hybrid and increasingly distributed environments, they are often challenged with balancing the complexity and rigidity of security with the flexibility and performance their applications require. Typically, while trying to balance these two competing forces, they err on the side of security, which can add rigidity. Still, when it comes to designing a security strategy focused on delivering the right protection across as many application architectures as necessary, flexibility is critical. As much as these two traits seem at odds with one another, the state of applications today requires them to coexist.

Organizations today need application-security strategies, technologies, and architectures that enable this diversity. They need the ability to shield their applications with the right protection, regardless of deployment or architecture circumstances. At the same time, these solutions should not impede modern development cycles, user experience, or necessary performance. Equally important is the need to deploy these protections with the same flexibility and speed as the apps they protect. There shouldn't be any tradeoffs, and with F5 Distributed Cloud Services, there doesn't have to be.

Distributed Cloud Services delivers capabilities to easily extend application and security services across one or more public clouds, hybrid and multi-cloud deployments, native Kubernetes environments, and edge sites. Distributed Cloud Services are differentiated in providing connectivity and security at both the network and application layers. As an overlay across separate cloud provider offerings, Distributed Cloud Services lets organizations easily integrate network operations, application performance optimization, security, troubleshooting, and visibility through a single management console. This delivers a platform-based approach that is cloud-agnostic and purpose-built to meet the needs of traditional and modern apps, all without increasing complexity or losing granular control and necessary visibility.

Secure Centrally, Operate Globally with Ease, with Distributed Cloud WAAP

WITH DISTRIBUTED CLOUD WAAP, ORGANIZATIONS CAN SECURELY CONNECT, DEPLOY, AND RUN APPS WITH CENTRALIZED MANAGEMENT, CONSISTENT SECURITY, AND END-TO-END OBSERVABILITY, ACROSS ANY INFRASTRUCTURE.

A major component of Distributed Cloud Services is Distributed Cloud WAAP, which protects and secures organizations' traditional, modern and hybrid apps with unparalleled performance and global scale. Distributed Cloud WAAP leverages a diverse set of security services with machine learning and globally-sourced F5 threat intelligence. These are operated across the F5 global delivery network, enabling SaaS-based application protection, including Web Application Firewall (WAF), API Discovery and Security, Bot Defense, and DDoS Mitigation.

With Distributed Cloud WAAP, organizations can securely connect, deploy, and run apps with centralized management, consistent security, and end-to-end observability, across any infrastructure or set of infrastructures. Apps can be made increasingly available to distributed audiences with a common set of security controls. Appropriate policies can be enforced wherever necessary, including within the F5 global network, within and across public/private clouds, and/or in an organization's data center in any combination. These are all deployed efficiently through a common user interface (UI) and central control plane. This document will help organizations decide what deployment architecture will suit their apps and existing infrastructure when implementing Distributed Cloud WAAP to protect their apps.

Distributed Cloud WAAP Deployment Options

SaaS Edge through Service Provider

The first deployment architecture supported for securing apps and APIs with Distributed Cloud WAAP is a traditional SaaS, service provider approach using a proxy to control the flow of application and API traffic to and from clients on the internet via the F5 global network. There are many reasons for a proxy to be deployed between clients and application origins, including layer 7 routing/load balancing, content caching (CDN), TCP optimization and, in

this case, applying critical application security policies to inspect, control and, if necessary, block traffic. This protects and secures against vulnerabilities, attacks, and other abuse and exploitation attempts.

In this scenario all applications requiring protection are advertised on the internet with a public IP or FQDN through “virtual host” proxies from the F5 global network. They can be protected and advertised globally through all Distributed Cloud Regional Edge (RE) ingress/egress and service delivery points of presence via anycast with distributed proxy architecture.

This is a very common and trusted approach to handling app and API security for distributed applications. By applying application and API security, including WAF, layer 3-7 DDoS Mitigation, and Bot Defense for web-facing assets at the F5 global network edge, we can stop attacks closer to their origin and help limit impact of bad application and API traffic before it hits a customer’s infrastructure. This can be a major benefit to overall performance and help to keep infrastructure and bandwidth costs down.




Figure 1: Clients connect to the closest F5 Global Network RE; traffic is targeted to a determined load balancer configuration, and security services are applied

SaaS Hybrid: Customer Edge (CE)

The second architecture supported for securing apps and APIs with Distributed Cloud WAAP is a local Customer Edge (CE) deployment. The uniqueness of the F5 Distributed Cloud SaaS platform and technology allows us to deploy local infrastructure workloads (software (VM) or hardware) in public or private clouds, all centrally managed and controlled via the global Distributed Cloud Services control plane. This allows us to push critical app and API security virtually anywhere there is computer, network, and storage available. Organizations can easily apply security controls, manage policy, and observe applications and APIs locally, whether in and across public cloud environments (AWS, Azure, GCP) or private cloud (on premises, in a data center or at the edge) where performance of an application is of the upmost importance, and applications don't require caching via CDN. This allows client requests and access to flow directly to a specific origin, for the most efficient routing and processing.

In this scenario, client requests can be directed to bypass the F5 global network and connect directly to the closest CE based on reachability design (GSLB, DNS resolution), wherever the app endpoint is deployed with localized, integrated layer 3-7 networking and security stack, which includes WAF, layer 7 DoS, rate limiting, and API discovery and protection. This deployment option can also assist in securing internal-only workloads that are not publicly accessible.

The value here is in the performance and experience for clients. Requests are allowed to go direct while remaining secure (privatized origin via mesh) so organizations don't have to sacrifice when it comes to the performance and security tradeoff. Distributed Cloud WAAP will streamline deployment, management, and observability for organizations with distributed apps and the need for direct, secure localized access as centralized management is delivered via the SaaS console. This can, however, introduce some complexity and potential cost impacts compared to a proxy-only architecture, including the need for organizations to manage advertisement of endpoints based on individually deployed apps and locations, and it doesn't allow for built in layer 3-4 DDoS Mitigation, via the Distributed Cloud Services global network. In this design, infrastructure costs could be significantly higher as requests are processed locally, even for all automated threats, Denial of Service (DoS), and other illegitimate traffic. This means organizations should expect more bandwidth is necessary for their local infrastructure, and they may incur greater hardware, maintenance, and management costs to maintain infrastructure necessary to support local Distributed Cloud VM nodes, along with their other app infrastructure.




Figure 2: Clients connect to the closest CE based on reachability design, including GSLB and DNS resolution

SaaS Hybrid: Regional Edge Plus Customer Edge

The third architecture supported for securing apps and APIs with Distributed Cloud WAAP is a hybrid app and API security deployment. This means a combination of proxied, publicly-advertised app and API endpoints with data paths through the F5 global network REs and VM CE node(s). This configuration allows organizations to provide security for applications they do not want directly exposed on the internet through a private subnet within a cloud services provider (CSP) or private data center, while other apps are easily advertised and protected everywhere in the world via the F5 global network.

Client access and requests to apps is granted via the closest RE of the F5 global network, advertised via Anycast for external, internet access. Incoming traffic, targeted to a determined Load Balancer, is routed appropriately, and security policies are applied with requests being passed upstream over the F5 global, private network through peered connections with CE nodes to securely access protected origin resources. In this scenario, the origin




Figure 3: Clients connect to the closest F5 Global Network RE based on anycast reachability; origin/destination is accessed via CE node or cluster through established tunnels to the local site FQDN/address.

is only accessible (internal DNS, private IP space) via CE node instances and the organization's supporting public or private routing environment. This removes the need for organizations to manage public DNS, public NAT, and firewall rules at their origin.

With this approach, organizations are best able to align application and API security with each application's needs, including optimizing architecture and, visibility, and control over their apps and APIs across environments. Naturally, this design merges some of the benefits and limitations of each, but the SaaS hybrid approach allows organizations to leverage a variety of application networking and security use cases that fit their organizations demands.

Leveraging CE VM nodes locally, in the cloud, across clouds, and within an organization's data center, provides capabilities to set up a flexible, scalable multi-cloud networking and security mesh supporting many modern app and multi-cloud use cases. This includes the ability to create virtualized Kubernetes clusters that span multiple clouds and CSPs to host and manage microservices workloads, a consistent network and security service platform with consistent layer 3 to layer 7 services, including virtual router, network firewall, distributed load balancer, WAF, layer 7 DoS, and API security. These can be deployed anywhere, enabling secure connectivity for highly reliable hybrid-cloud connectivity. Further connecting these workloads and environments streamlines multi-cloud app management and delivers a unified end-to-end policy with granular observability across clouds and distributed workloads.

In this architecture, organizations can also create a second tier of localized app and API security enforcement, allowing secure, direct user/client access to critical endpoints where and when necessary. The hybrid design with Distributed Cloud Services also allows organizations to deliver a more effective, defense-in-depth, layered application security approach, even when security requirements vary by app or environment. Organizations are easily able to apply baseline app security for all applications through the F5 global network REs (such as layer 3-7 DDoS Mitigation, and general WAF policy), while delivering capabilities for localized app or environment specific app and API security policies to be applied within CSP, data center, or edge environments. This also creates support for critical app-to-app connectivity and security, east/west between environments and apps, while delivering and securing all external client connections north/south for any app.




Figure 4: Logical flow of client app requests through F5 Distributed Cloud WAAP security functions

Why the Flexibility of Distributed Cloud WAAP Matters to Your Organization

Organizations need application security solutions that deliver, secure, and help them maintain compliance in this increasingly complex hybrid app world, while not hindering innovation and performance. It's critical that organizations match app and API security with deployment and architectural specifications. They need security solutions that support the needs of hybrid applications, that move at the speed of applications. That's exactly what Distributed Cloud WAAP can do, delivering app security:

- At scale, to keep up with increasing attack sizes/scope and automation, but also be able to expand to support security for an increasing number of apps and environments
- With the flexibility to support apps anywhere, working within any environment, without requiring changes to architecture approach or optimal data paths
- Which drives simplicity, requiring minimal CAPEX, streamlining operations while providing necessary visibility and control of apps, no matter where they need to be deployed

Distributed Cloud WAAP and the Distributed Cloud Platform provides a single SaaS-based control plane and set of services to deploy, manage, secure, and observe applications and APIs anywhere. Each deployment model can help organizations simplify security and improve visibility while reducing operational complexity. Whether delivering globally or locally, the solution provides the security and services to meet customers' application security needs as they deliver and scale wherever their applications are hosted.

Summary of the Functional Considerations of the Various F5 Distributed Cloud WAAP Deployment Options

	RE-Only Deployment	CE-Only Deployment	Hybrid Deployment
Primary Traffic Proxy (where is the proxy?)	F5 Regional Edge (RE)	Customer Edge (CE)	F5 Regional Edge (RE)
WAF	Yes	Yes	Yes
API Security	Yes	Yes	Yes
Bot Defense	Yes	Yes	Yes
Layer 3-4 Routed DDoS Mitigation	Yes	No (requires separate subscription)	Yes
Layer 7 DDoS Mitigation	Yes	Yes	Yes
Centralized Management and Visibility	Yes, with Distributed Cloud Console	Yes, with Distributed Cloud Console	Yes, with Distributed Cloud Console
Load Balancing	Yes	Yes	Yes
Delegated DNS	Yes	No	Yes
Anycast	Yes	No	Yes
Private Origin	No, publicly accessible origin (public FQDN, NAT)	Yes	Yes
Requires Hosted Component	No	Yes	Yes

Check out the F5 Hybrid Security Architectures Series on [F5 DevCentral](#). To learn more about how Distributed Cloud WAAP helps secure your apps and APIs, visit f5.com/waap.

