F5 Application Delivery and Security Services on AWS

KEY BENEFITS

- Industry-leading protection for applications and data on AWS
- Federated access for any user, on any device, anywhere within your multi-cloud architecture
- Maintain and ensure high availability of web applications for enhanced user experience
- Workload portability between AWS, other cloud platforms and on-premises deployments
- Increase agility of deployments through automation
- Improve control and manageability with a single set of highly programmable application services

aws partner network

Advanced Technology Partner

Networking Competency

Security Competency

F5 and Amazon Web Services (AWS) have partnered to help you rapidly deploy and secure application workloads without incurring the capital expenditures of new infrastructure. A pioneer and leader in the public cloud market, AWS provides customers with the ability to provision and deploy application workloads in the AWS cloud—and only pay for the application and infrastructure resources as needed. F5 takes you beyond the basic load balancing, networking, and security services available in AWS by delivering advanced and programmable L4–L7 application availability, access and security services consistent with those available in data center or other multi-cloud deployments.

Challenge

Infinite scalability, unmatched flexibility, reduced overheads—the benefits of migrating or developing apps in the public cloud are obvious. However, many enterprises making the shift to the cloud do so amid concerns that their applications' security and performance may be diminished, or that they'll fall victim to cloud vendor lock-in and incur large re-architecting costs should they need to relocate workloads. Increased network segregation and inconsistent application services across multi-cloud architectures are placing additional strain on IT departments, while creating loopholes that can result in new security vulnerabilities.

But it doesn't need to be this way.

Solution

F5® BIG-IP® Virtual Editions (VEs) have been fully tested and verified by F5 for use in the AWS cloud, and are recognized and endorsed by AWS through attainment of both their Security and Networking Competencies. Because all BIG-IP devices are built using the same base code, Virtual Editions offer complete feature and functionality parity with BIG-IP hardware devices. As a result, by employing F5's advanced L4-L7 traffic management and security services to front your AWS applications, you won't have to worry about reduced security posture, application performance, or availability. Not only that: Policies and configurations from pre-existing BIG-IP deployments can be leveraged and replicated within AWS, reducing time to

market. Deployment times can be shortened further through use of F5's extensive range of CloudFormation Templates, enabling fast, reliable and automated VE deployments. A number of highly flexible licensing models are available for VEs on AWS, including Bring-Your-Own-License (BYOL) and Pay-as-You-Go (PAYG) consumption via the marketplace, or alternatively F5's annual subscription and enterprise licensing agreement (ELA) are also both supported.

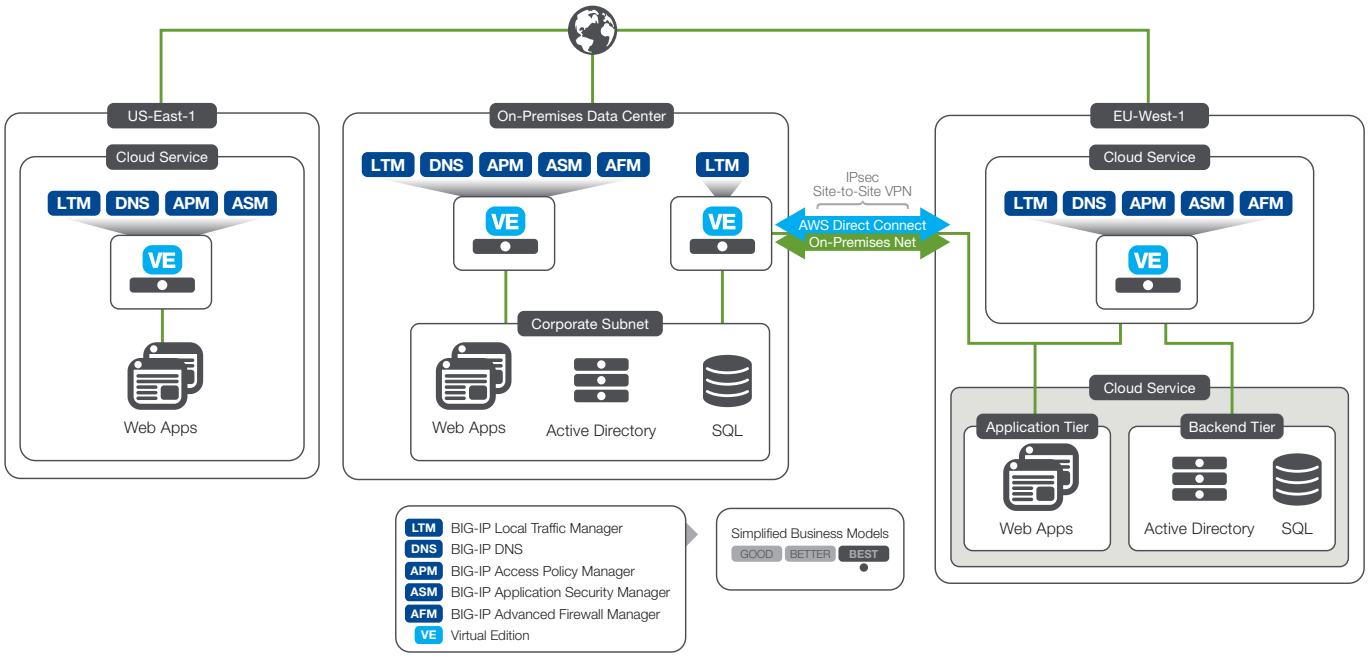

Protect AWS Workloads with Robust, Enterprise Class Security Services

All public cloud providers take the security of their platforms incredibly seriously. However, there are two distinct categories of security in the public cloud: security *of* the cloud and security *in* the cloud. Security of the cloud is the cloud provider's responsibility, and entails anything relating to the security of the underlying platform infrastructure—including compute resources, databases, networking and physical data center security. In contrast, security in the cloud is the sole responsibility of the application owner, and relates to the security of individual applications and their data within the cloud platform.

Running BIG-IP VEs in your AWS environment and implementing F5's advanced L4-7 security services is the easiest and most effective way of ensuring your applications, network, and data are protected from cybercriminals in the cloud. At the network level (L4), the BIG-IP Advanced Firewall Manager (AFM) scales to shut down high capacity DDoS attacks that can overwhelm load balancers, firewalls and complete networks. At the application layer (L7), BIG-IP Application Security Manager (ASM)—F5's industry-leading Web Application Firewall (WAF)—mitigates against common application vulnerabilities and L7 DDoS attacks, while providing protection against all OWASP top 10 threats. Should you already employ either of these modules anywhere else in your multi-cloud world, you can quickly and easily replicate security policies in AWS to ensure consistent security across your architecture, preventing loopholes which may arise through use of multiple security products.

KEY FEATURES

- Consistent application services across AWS and on-premises deployments
- Protection from bots, L4–L7 DDoS attacks, and all OWASP top 10 threats
- Secure, policy-driven single sign-on (SSO) and federated access
- SSL offloading and stateful L4–L7 traffic management
- Tools for automation and programmability

Complement AWS Native WAF with F5 Managed Rules

Alternatively, should you establish that some of the workloads in your application portfolio don't require the advanced security offered by F5's WAF, but feel the protection offered by AWS' native WAF is inadequate, then F5 has partnered with AWS to deliver sets of managed rules that can be implemented on top of the AWS WAF to provide an extra layer of protection. These pre-configured rulesets are written and managed by F5 security experts, updated on a regular basis, and cover a range of security threats including bot protection, common vulnerabilities and exposures, and web exploits that are included in the OWASP top 10. These rules can be added any time with a few of clicks, and are available for consumption on a pay-as-you-go basis that ensures you only pay for what you use.

Increase Flexibility and Scalability with Cloud Bursting

The promise of limitless scalability causes many to move applications entirely to the public cloud. For those who are still slightly wary of the cloud, this scalability can be exploited in another way: cloud bursting. This deployment model allows an application to run primarily within a data center or private cloud environment. When traffic exceeds capacity, any additional traffic is directed to, and absorbed by, identical instances running within AWS. Designing a federated cloud in this way has many benefits from an economic standpoint: deploying BIG-IP VE's into this arrangement enables fast, seamless, geolocation-based redirection of application users over secure SSL-VPN connections. The user experience remains unaffected regardless of whether the responding server is located on-premises, or on AWS.

Improve Performance and Availability with Global Traffic Management

The ability to replicate applications throughput multiple geographic region across AWS and on-premises data centers empowers application owners to improve redundancy. It also reduces the physical distance between an endpoint device and an application server, providing lower-latency access to device users. Implementing BIG-IP DNS VE in your cloud network lets you go one step further by using global server load balancing to make informed routing decisions based on physical proximity of a server, the real-time performance and health of a server, or a number of other metrics to ensure end users have the best possible experience.

Provide Federated Access to Your AWS Network and Applications

Installing BIG-IP VEs into your multi-cloud environments solves the problem of federating access, network, and application resources across your data center and AWS environments. BIG-IP Access Policy Manager (APM) uses Security Assertion Markup Language (SAML) to enable web browser Single-Sign-On (SSO), multi factor authentication, geolocation restricted access, and device inspection. SAML also eliminates the need to manage independent user accounts across Software-as-a-Service (SaaS) providers. BIG-IP APM simplifies and secures access management, with identity, context and application aware policies, enabling access from all networks and devices

Improve Efficiency of AWS apps with Advanced Programmability

F5 iControl is an open web-based API that provides complete dynamic control of F5 configuration objects. You'll have the power and flexibility to ensure that applications and their underpinning network—whether in AWS or on-premises—work together efficiently to simplify management of complex architectures. In addition, you can use F5s iRules scripting language to provide complete programmatic access to traffic flowing between to and from applications. iRules allows you to inspect, analyze and redirect traffic entirely based on your custom ruleset.

Boost Deployment Agility in AWS with Automation Tools

Deploying applications in the cloud should always be a fast, effortless process. However, this is only achievable if the supporting application services can be fabricated in a similar fashion. With F5 designed and tested AWS CloudFormation templates, BIG-IP VEs can be spun up in a matter of minutes across a wide range of topologies and use cases—including anything from standalone VEs to completely autonomous autoscaling solutions. Using these templates, located in F5's GitHub repository, everything from the deployment of essential AWS resources to the configuration of the BIG-IP VE is performed in just a few clicks. Not only that, a number of these solution templates have been fully integrated into the AWS marketplace so they can be deployed directly into your AWS VPC from the marketplace—further reducing time to market.

Alternatively, F5 iApps templates can rapidly configure BIG-IP VEs to best suit the requirements of a specific application, based on a few simple checkbox inputs. These can then be re-used to configure any BIG-IP device and replicate the configuration settings across a multi-cloud architecture. In this way, iApps Templates reduce IT time consumption and ensure policy consistency across your deployment.

Enjoy Flexible Licensing and Consumption Models

To better align with public cloud-based usage models, BIG-IP VE offers four distinct consumption options that give you the flexibility to meet specific operational and financial requirements.

- Pay-As-You-Go (PAYG) is available for those looking to leverage F5 ADC services on a per-hour basis; this option is perfect for development and test environments, or short-term projects in AWS.
- F5s Subscription Licensing lets you purchase one, two, or three-year BIG-IP VE license subscriptions that can be deployed in any supported environment. You can self-license additional BIG-IP VE licenses using BIG-IQ License Manager.
- F5s Enterprise Licensing Agreement (ELA) provides you with the architectural flexibility to deploy however many VEs you want, wherever you want, and whenever you want, with no retroactive penalties that can wreak havoc on budgets.
- F5s Bring-Your-Own-License (BYOL) option is a perpetual license that lets you amortize acquisition costs of a longer period of time.

Conclusion

The process of migrating to, or developing on, AWS can be greatly simplified and accelerated with F5's application delivery services, while dramatically increasing security, performance and availability of applications. F5 provides an abstracted tier of application services that can be implemented consistently, anywhere within a multi-cloud architecture, eliminating the need for multiple disparate solutions and the resulting IT strain. This deployment approach enables enterprises to seamlessly and confidently extend private data centers into the cloud.

DC0318 | OV-CLOUD 209914596 | 0318