



Migrating Tier 1 Application Workloads to AWS with F5

Enterprises deploying applications on AWS may have concerns about security, performance, and loss of control. The F5 BIG-IP platform alleviates these concerns by providing the same high availability, performance, and security services that most enterprises currently have in their private data centers.

White Paper
by F5



WHITE PAPER

Migrating Tier 1 Application Workloads to AWS with F5

Introduction

Enterprises that are considering migrating or deploying production workloads in the public cloud are most likely to adopt Amazon Web Services (AWS) as their preferred platform. According to Gartner, AWS is the overwhelming market share leader, with over 10 times more cloud Infrastructure as a Service (IaaS) compute capacity in use than all 14 other providers combined in the same magic quadrant¹. As the public cloud platform service of choice, AWS supports the broadest range of use cases, including those for enterprise and mission-critical applications². This paper reviews many of the factors that have inhibited broader adoption of the public cloud for Tier 1 applications and suggest how F5 application delivery services play a critical role in helping accelerate AWS adoption.

The Need to Rapidly Deploy Infrastructure and Services

For most enterprises, Tier 1 applications are not static resources with clearly defined performance parameters. Frequently, enterprises must adapt to unexpected spikes in demand for a particular application. Traditionally, IT organizations have responded to this demand by provisioning infrastructure and application resources to meet peak rather than average demand. This type of resource overprovisioning was necessary because of the inflexibility of in-house IT services and the need to invest in fixed assets. But procuring and deploying new infrastructure, networking, and associated application resources is a complex and time-consuming process that requires considerable CapEx investment.

AWS potentially solves this part of the challenge for enterprises because it has already made the necessary investments in infrastructure and software assets, thus enabling enterprises to reap immediate benefits. The advantages of using a public cloud IaaS provider like AWS is the agility it affords enterprises to rapidly deploy new applications, the flexibility to allocate resources on demand, and the economics of an OpEx versus CapEx model.

Enterprise Concerns about Moving Applications to the Public Cloud

It should be no surprise that application security, performance, and management control are among IT's top concerns when they consider moving their tier 1 applications to the cloud. While AWS provides many native tools and services to address some of these application delivery factors, those tools have varying capabilities and feature sets that might not meet your application requirements.

Security



WHITE PAPER

Migrating Tier 1 Application Workloads to AWS with F5

Given today's threat landscape, most organizations rank security as their number one concern for cloud-hosted applications. It's critical to protect against the sophisticated malware and blended layer 4–7 security threats, such as volumetric DDoS combined with application layer attacks (OWASP Top Ten, cross-site scripting, SQL injection, etc.) when moving workloads to any public cloud infrastructure. AWS uses a shared responsibility model with regard to security; you are responsible for securing everything above the hypervisor layer, namely your applications.

Consistent application access is also a requirement. How do you ensure that access is uniform and conforms to your organization's policies? Password fatigue can compromise security if users have to memorize multiple passwords. So, you also must be able to replicate and enforce consistent and proven security practices and policies across your private and public cloud environments.

Performance

User experience and productivity continue to be important considerations, and both are dependent on how well applications perform once in the cloud. In some cases, the cloud provider's data center will be farther away from your users, which means there will be increased latency between user and application, impacting performance. In addition, some of the methods that are typically used to deal with latency, such as caching, compression, and TCP optimizations, are not available in AWS.

Advanced Traffic Management and Visibility

Most enterprises require advanced traffic management (beyond basic load balancing) in their data centers for their business-critical applications. While AWS offers basic load balancing services via elastic load balancing (ELB), you should consider what protocol support beyond HTTP/HTTPS and TCP you will need. Will basic health checks and load-balancing algorithms be sufficient? Often, you need to be able to manipulate application data, which requires full L7 application proxy functions, such as URL inspection and rewriting. The ability to see incoming client traffic with context is critical to your ability to make granular traffic steering decisions.

How F5 Enables Adoption of AWS

Addressing the above concerns requires advanced and programmable application delivery and security services on a unified platform like F5 BIG-IP. This platform ensures the security, performance, and availability of all applications, regardless of their location. It also enables you to deliver and manage application services and associated policies in a consistent way across all application environments for both existing applications and new cloud-based applications.



WHITE PAPER

Migrating Tier 1 Application Workloads to AWS with F5

F5 BIG-IP virtual editions (VEs) are virtual appliances that deliver the same set of services that are available on F5 BIG-IP hardware devices—intelligent application and networking services, from acceleration, optimization, and intelligent traffic management (both local and global), to DNS, advanced application access, and application security. These services can be fully integrated as part of the application stack and configured automatically. As the market leader in both hardware and virtual Application Delivery Controllers (ADC), it's possible that F5 services are already deployed in your data center.

F5 Security Services

BIG-IP VEs deliver intelligent, comprehensive L4–7 security services that protect cloud applications without requiring you to sacrifice control, flexibility, and visibility. These services complement AWS offerings and provide defense in depth against the full spectrum of DDoS vectors, web scraping, multilayer web-based application attacks, data theft, and leakage. With the intelligence and advanced behavioral analysis to recognize anomalous traffic patterns, F5 solutions can detect and mitigate automated botnet attacks. In addition, with the power of F5 iRules® data path scripting, F5 solutions can quickly respond to exploits of application vulnerabilities and zero-day attacks. With F5, the effort and expertise you invest in tuning and configuring the firewall rules and policies for each in-house application can be leveraged and reused in the cloud.

F5's identity and access management architectures are based on full user, device, environment, application, and network context awareness. That means F5 solutions enable identity federation and single sign-on for application access across the data center and the cloud. At the same time, they enable you to maintain the security of the applications and integrity of data with secure, differentiated access based on context, protection from web-based malware and persistent threats, and comprehensive endpoint device inspections.

Availability and Performance

BIG-IP advanced local traffic management services support a broad range of protocols beyond HTTP/TCP (for example, HTTP 2.0, SPDY, and UDP), and deep application fluency. As a full application proxy, the BIG-IP platform enables content switching in order to mitigate the limited number of external IP addresses provided by AWS. It also dynamically tracks the performance levels of servers in a group and provides deep health monitoring and connection state management. BIG-IP application delivery optimization services can accelerate your application response time, minimize latency and delays, and reduce the number of data round trips necessary to complete web requests from mobile devices.



WHITE PAPER

Migrating Tier 1 Application Workloads to AWS with F5

BIG-IP DNS and global server load balancing services direct users to the nearest cloud data center based on the best application experience and disaster recovery and failover policies. User proximity, geolocation, network conditions, and application availability are all factored into routing decisions. The platform employs a range of global load balancing methods and intelligent monitoring specific to each application and user. F5 also offers DNS DDoS protection, blocks access to malicious IP addresses, and secures responses with DNSSEC. Best of all, DNS queries and health checks are not billed per use, so you avoid the high cost of being charged for both legitimate and illegitimate queries during a DNS DDoS attack.

Scalability to Meet Application Performance Requirements

A key benefit in moving application workloads to public cloud platforms is the ability to scale an application beyond your provisioned base capacity in the data center. With AWS Auto Scaling (see Figure 1), you can maintain application availability while automatically scaling your Amazon EC2 capacity up or down according to pre-defined thresholds. BIG-IP solutions integrate with AWS Auto Scaling to enable you to dynamically scale BIG-IP application and security services. In addition, because BIG-IP VEs natively handle the addition or removal of pool members, there is no need for out-of-band orchestration and configuration management.

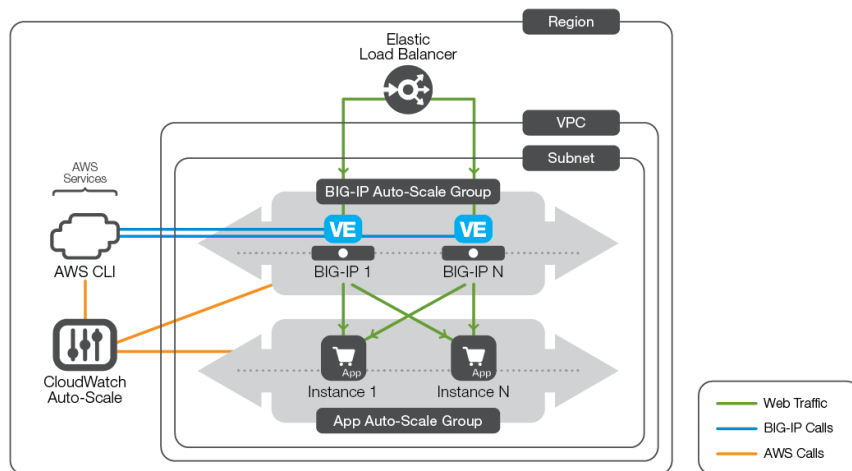


Figure 1. Automatically scale F5 application and security services along with Amazon EC2 instances.

Simplified Deployment



WHITE PAPER

Migrating Tier 1 Application Workloads to AWS with F5

AWS Cloud Formation Templates (CFTs) provide a scripted method for automating the deployment of infrastructure (server, storage, networking, and compute) resources. They provide a repeatable way to rapidly deploy the same configuration multiple times within AWS.

To simplify deployment of F5 application services in AWS, F5 provides sample CFTs on the F5 DevCentral™ community site. There, you'll find example configurations of VPC resources such as subnets, network interfaces, and routing tables for deployments of BIG-IP VEs. These examples also show how you can use CloudInit user data scripts to deploy BIG-IP iApps® templates for specific packaged applications (Microsoft SharePoint, Exchange, and others) and custom applications. Similar in functionality of AWS CFTs, F5 iApps were created to help you quickly deploy the specific services each of your applications needs. You can also use iApps to define the configuration and policies of services such as traffic management, encryption, firewall, and performance optimization for each application.

Federated Public and Private Deployments

Integrating public cloud resources with your existing private data center enables you to transition your application workloads based on prioritized schedules while continuing to leverage your existing investments. Using BIG-IP VE's, F5 iApps, and AWS CFTs together, you can create an integrated cloud configuration that lets you rapidly and transparently deploy additional application resources. Key advantages to this federated cloud configuration include the seamless redirection of application users, geolocation and acceleration technologies, and secure connections using AWS Direct Connect. The user experience remains unaffected whether you choose to deliver application resources from a private data center or public cloud.

Transparent and continuous use of private and public cloud resources can be based on demand, on whether a project is new or already in place, or on the specific location of the requester.



WHITE PAPER

Migrating Tier 1 Application Workloads to AWS with F5

Integrating the management tools and connectivity between public and private environments creates a seamless experience across the two, delivering a transparent extension to the data center environment and breaking down management silos. The F5 BIG-IQ® management and orchestration platform provides an open, programmable, and intelligent framework for centrally managing physical and virtual BIG-IP solutions across both private and public cloud infrastructures (see Figure 2). BIG-IQ enables you to deploy and manage application and security services in a fast, consistent, and repeatable manner—regardless of the underlying infrastructure. In addition, BIG-IQ integrates with AWS thru REST APIs and specific AWS connectors. It also provides lifecycle management of iApps, which simplifies, automates, and accelerates your ability to provision application delivery services. Your IT organization can create a Service Catalog of available application delivery and security services, from which administrators and application developers can quickly select services as needed.

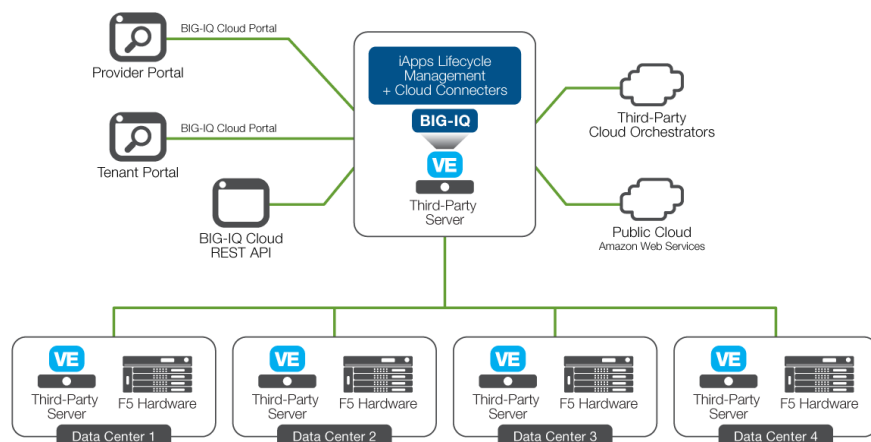


Figure 2: The BIG-IQ platform manages federated private and public cloud environments.

Flexible Licensing Models

BIG-IP virtual editions are available in Good-Better-Best bundles and in three distinct purchase models.

- Utility billing. For test and development pre-production workloads or temporary cloud bursting and scalability use cases, F5 offers hourly utility licensing that provides flexibility and on-demand use (pay for only what you use, after you use it). Annual subscriptions are ideal for production workloads that have steady-state traffic.
- Bring-your-own-license (BYOL). This option is ideal for hybrid cloud environments in which you want to take existing BIG-IP VE licenses from your private data center directly into AWS.
- Volume Licensing Subscription (VLS). This offering is designed for enterprises



WHITE PAPER

Migrating Tier 1 Application Workloads to AWS with F5

that require large volumes of virtual application and security services. VLS provides discounted pricing for 1- and 3-year subscriptions with premium support and software updates included.

Conclusion

Adoption of public cloud services has grown exponentially over the past few years, and AWS has been the consistent market leader for these services. Many IT startups and even some large, well-known enterprises have deployed entirely in the AWS cloud, achieving significant success. As enterprises plan to migrate more critical applications to the cloud, F5 application delivery and security services can be easily ported to cloud application workloads using F5 BIG-IP virtual editions. Doing so addresses many of the fundamental concerns that enterprise customers have regarding security, performance, and control in the public cloud. With flexible licensing models available for BIG-IP solutions in the AWS Marketplace, enterprises can plan, stage, and deploy applications in AWS with minimal financial exposure and begin to benefit from the agility and efficiencies of AWS cloud.

¹ [Gartner Magic Quadrant for Cloud Infrastructure as a Service May 18 2015](#)

² [Gartner Magic Quadrant for Cloud Infrastructure as a Service May 18 2015](#)

F5 Networks, Inc.
401 Elliott Avenue West, Seattle, WA 98119
888-882-4447 www.f5.com

Americas
info@f5.com

Asia-Pacific
apacinfo@f5.com

Europe/Middle-East/Africa
emeainfo@f5.com

Japan
f5j-info@f5.com