Migrating Application Workloads to Public Cloud

Organizations migrating applications to public IaaS providers must continue to deliver an outstanding end-user experience while maintaining security, visibility, and control. F5 application and security services can achieve these goals while providing agility, consistent application and security policies, and operational cost-efficiency.

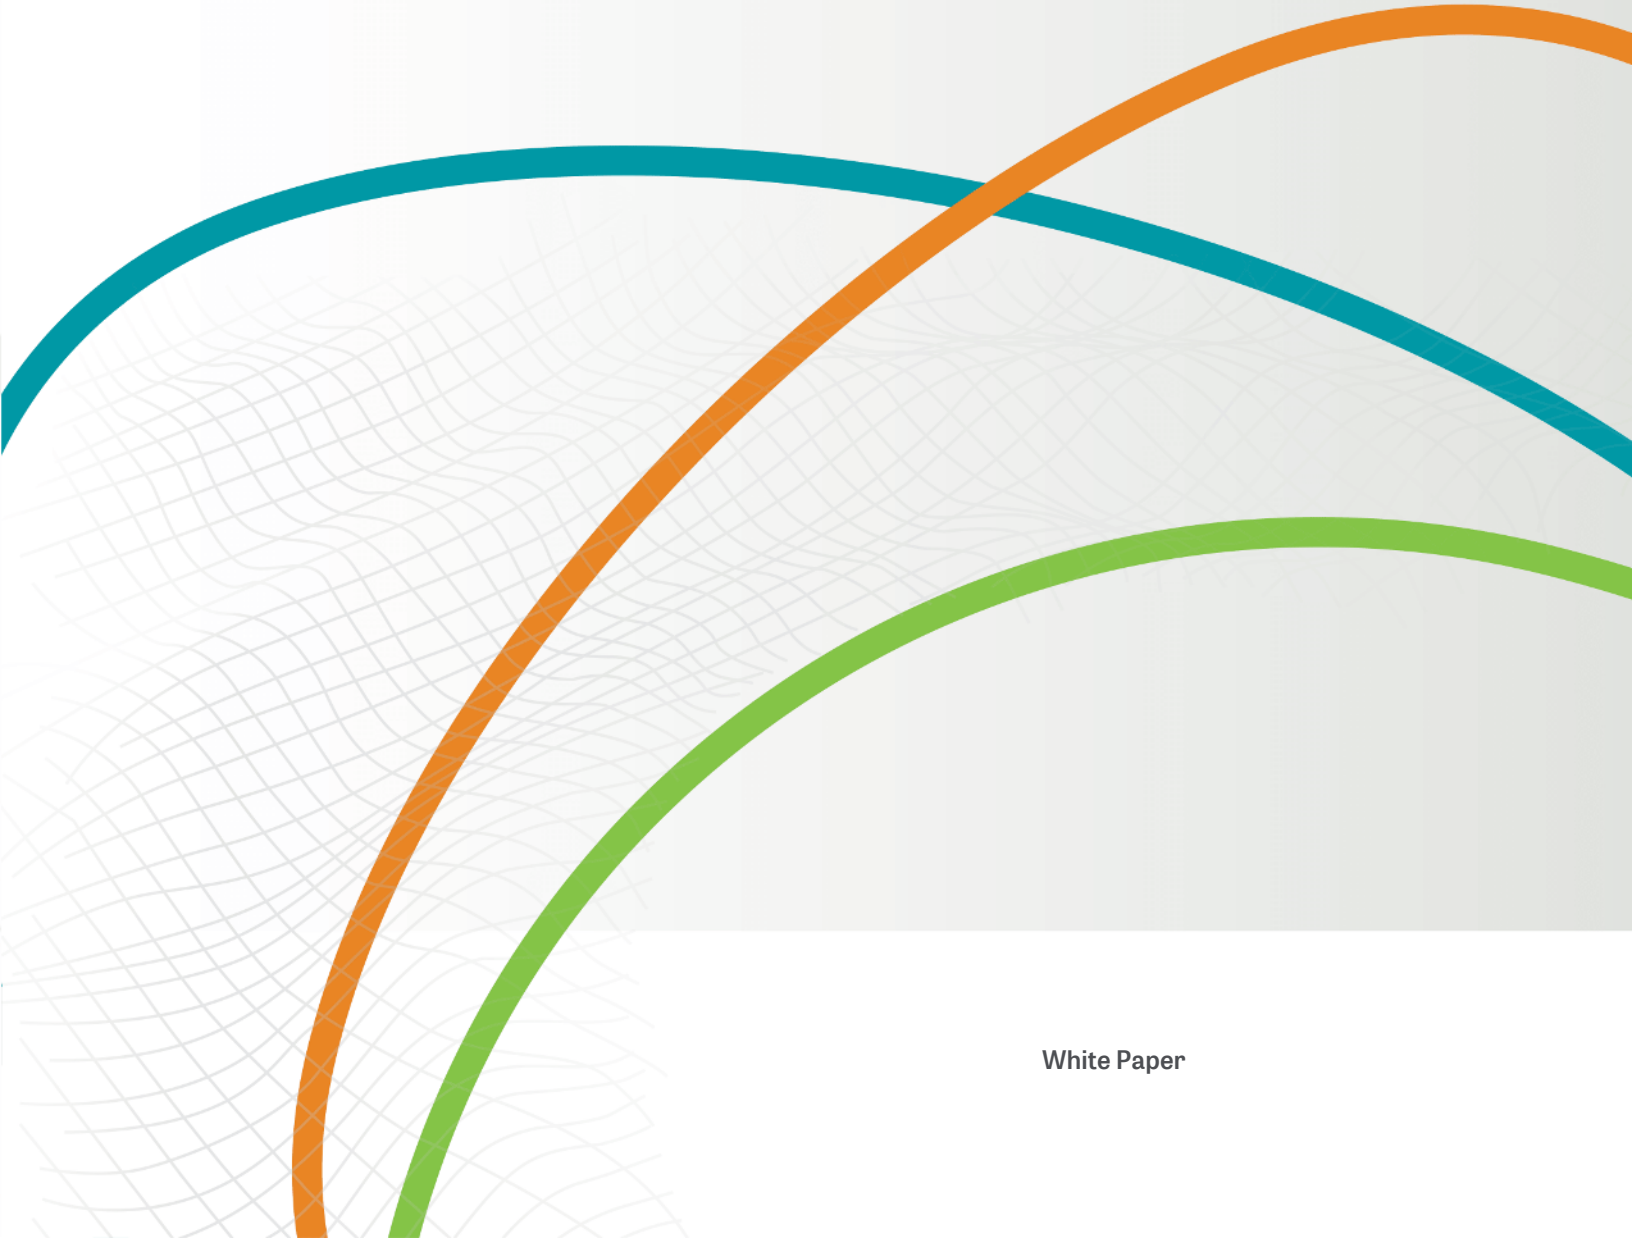

White Paper

Introduction

Enterprises and organizations from all industries and sectors are migrating or deploying new applications to IaaS public cloud providers to achieve greater agility, faster time to market, and flexible utility payment models. Whether these applications are revenue generating or critical business apps, they must ensure the same great user experience, including across associated availability, performance, and security services. However, there are challenges that need to be addressed, including determining which workloads are suitable for the cloud due to the inherent design of cloud data centers, the application delivery and security capabilities of each cloud provider, and the overall lack of visibility and control.

These challenges may lead to slow and expensive customized cloud implementations, prevent cloud vendor choice and mobility, and increase the risk of security vulnerabilities. This paper will review key considerations and known inhibitors to the successful migration of applications to the public cloud, and how deploying F5® BIG-IP® application and security services—available with flexible licensing models in the leading cloud IaaS providers—are a critical element of that success.

Cloud adoption barriers

Source: 2015 CLOUD SECURITY SPOTLIGHT REPORT, Information Security Community on LinkedIn

Challenges in Migrating Applications to the Public Cloud

While we all recognize the benefits of the public cloud, the fact is that there are significant differences between how an application runs in a public IaaS provider data center designed for multiple tenants and how it runs in your private enterprise datacenter. The public cloud provider will have designed its data centers and networks with massive scalability in mind, using virtualization, commoditization, and standardization to drive down costs. The level of network control is different, access to L2 functionality (e.g., multicast, 802.1q VLAN tagging, etc.) will be limited, and you may only get one public-facing IP for each application. Adding compute capacity is done by scaling out many small instances versus scaling up via high-performance dedicated hardware, which directly impacts maintaining state on any particular element or node.

Application workloads suitable to move to the cloud

Though the ideal goal is to "lift and shift" to the public cloud, some applications may not be able to be moved without design changes or being completely re-architected. Deciding which applications to migrate to the public cloud and what changes are needed is key. Questions and criteria for deciding which apps to move should include:

- Is the application virtualized already, and can it run on the cloud provider's infrastructure?
- Does the cloud provider meet your strict data protection policies?
- What level of service-level agreement is required for your application? What type of uptime guarantee or high-availability capabilities does the cloud provider offer?
- What are your network and management requirements? Can they be duplicated in the cloud environment?

Application and security concerns

Every application requires app and security services regardless of location. Each cloud provider's toolsets and services for availability, performance, and security will differ in capabilities and management, and may incur additional costs that need to be factored in. Learning and configuring these new services for your requirements will require time, testing, and training. This can create prohibitive switching costs when using a multiple-cloud-provider strategy, resulting in cloud vendor lock-in. Most important, depending on your specific application requirements, they may not be adequate nor provide the same capabilities as what you use today. This can limit business flexibility in choosing cloud providers, and increases the risk and complexity of cloud migration. There are three principle challenges:

1. Security

90% of organizations have security concerns

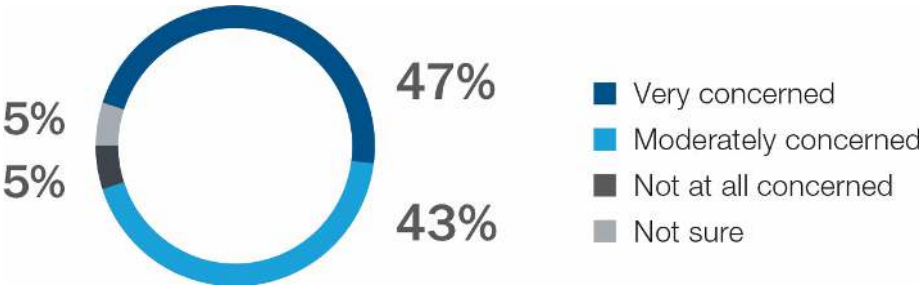

Source: 2015 CLOUD SECURITY SPOTLIGHT REPORT, Information Security Community on LinkedIn

Making sure that applications are secure in the public cloud is the top concern for most organizations. Protecting against the sophisticated, blended L3–7 security threats, where multiple types of volumetric DDoS attacks are combined with app layer attacks (OWASP Top Ten, cross-site scripting, SQL injection, etc.) is critical. Another consideration is the inconsistency of access and application security policies when using the cloud provider's basic security tools. This can increase attack surfaces and expose the vulnerabilities related to provisioning and de-provisioning access for users, especially the bad actors. Organizations need the ability to replicate and enforce consistent and proven security policies and access across the private data center and cloud.

2. Availability

Advanced traffic management beyond basic load balancing is typically deployed for business-critical and major enterprise applications. While cloud providers may offer basic load-balancing services, you should consider what protocol support beyond HTTP/HTTPS and TCP will be needed. Are basic health checks and load-balancing algorithms hash-based and round robin sufficient? Application data manipulation is often needed, which requires full L7 proxy functions, such as URL inspection/rewrite.

How does the cloud provider address uptime and resiliency of its infrastructure? Typical infrastructure availability targets are 99.95% for the larger providers, which may be lower than what you require. More important, understanding the risks and finding ways to mitigate the effect of an outage are critical. Some providers have redundant data centers and locations in multiple geographic regions to maintain availability in the face of major failure modes, such as natural disasters. Leveraging the redundancy requires careful planning that factors in specific implementation details, latency, and failover/recovery times.

3. Performance and scalability

End-user experience and productivity will continue to be vital and are dependent on how well the application performs once in the cloud. The data center may be farther away from your users, which means increased latency between the end user and application, impacting performance. Some of the methods that are typically used, such as caching, compression, and TCP optimizations, may not be available. Ensuring that users get directed to the closest location is another requirement.

One of the key reasons for going to a public cloud is to gain flexible, on-demand allocation of resources to address spikes in demand—planned and unplanned—based on predefined thresholds. Applications where a short-lived burst in capacity or highly variable demand can be expected may be better candidates for temporary migration to a public cloud provider.

Lack of visibility and inconsistent policies

Application and security services and policies that don't follow applications from the data center to the cloud will require customized implementations for each cloud provider. Organizations need insight into application performance, security, and application access by users to determine when workloads should be moved from one location to another. This requires visibility into user interactions with applications and the user experience across all deployment infrastructures. Policy sprawl, variability, and complexity per app and per provider, combined with lack of coherent visibility, can lead to increased OPEX, reduced service velocity, and a degraded customer experience.

How F5 Enables Successful Cloud Migration

BIG-IP platform for delivering application and security services

Addressing the above challenges requires advanced and programmable application delivery services that can span private data centers and cloud providers, providing business flexibility and enabling a successful cloud migration. Organizations need a unified platform that enables them to deliver and manage application services and associated policies in a consistent way across their application environments for existing applications as well as new cloud native applications.

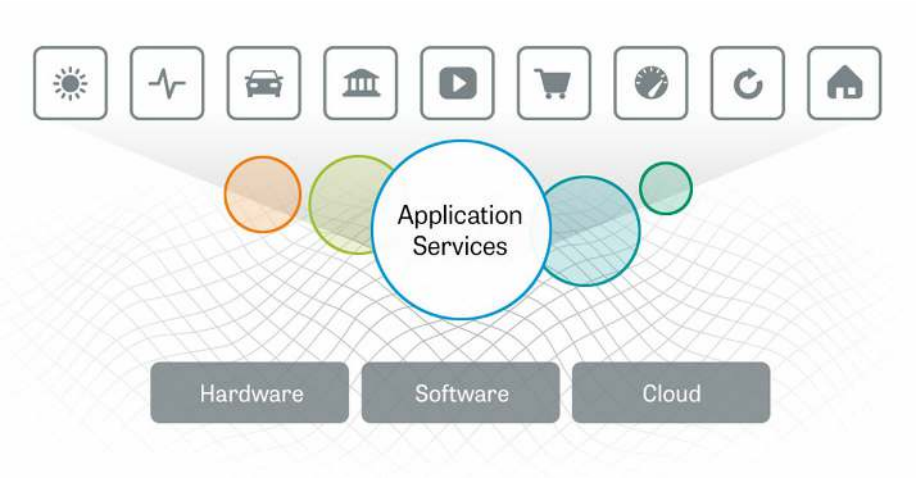

F5 application services on a unified platform for the cloud

F5 BIG-IP virtual editions (VEs) deliver a broad range of intelligent application and networking services, from acceleration, optimization, and intelligent traffic management (both local and global), to DNS, advanced application access, and network 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), these services are likely to be already deployed in your data center and are now available from leading cloud providers through F5's cloud licensing program.

F5 Security Services

BIG-IP VEs deliver intelligent, comprehensive L3–7 security services that protect cloud apps without sacrificing control, flexibility, and visibility. These services complement what each cloud provider offers and provide in-depth defenses against the full spectrum of DDoS vectors, zero-day threats, multilayer web-based application attacks, data theft, and leakage. The effort and expertise involved in tuning and configuring the firewall rules and policies for each application can be leveraged and reused in the cloud provider.

With identity and access management architectures based on full user, device, environment, application, and network context awareness, F5 enables identity federation and single sign-on for application access across the data center and into the cloud—all the while maintaining 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.

Consistent security across IT infrastructures and continuous protection are the most important factors for protecting cloud environments.

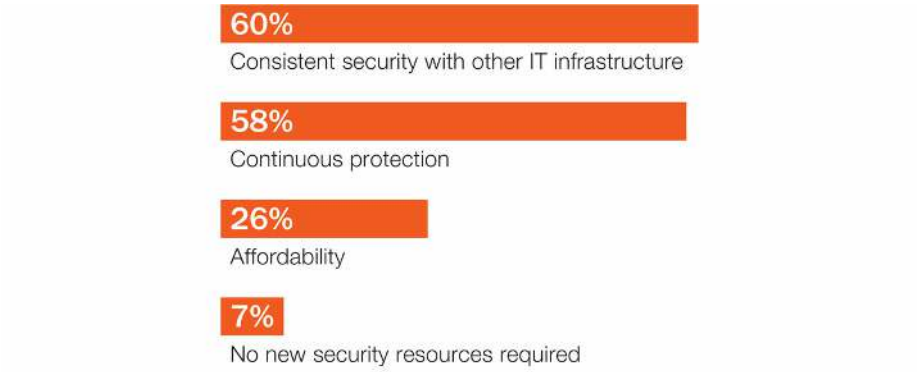

Source: 2015 CLOUD SECURITY SPOTLIGHT REPORT, Information Security Community on LinkedIn

Availability and performance

BIG-IP advanced local traffic management services support static and dynamic load balancing to eliminate single points of failure, a broad range of protocols beyond HTTP/TCP (e.g., HTTP2.0, SPDY, UDP), and deep application fluency. As a full application proxy, the BIG-IP platform enables content switching for multiplexing to mitigate a limited number of external IPs. It also tracks the dynamic 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.

The BIG-IP platform with DNS and global server load-balancing services directs users to the nearest cloud data center based on best application experience and DR/failover policies factoring in user proximity, geolocation, network conditions, and application availability. The platform employs a range of global load-balancing methods and intelligent monitoring specific to each application and user. F5 also provides DNS DDoS protection, blocks access to malicious IPs, and secures responses with DNSSEC. Best of all, DNS queries and health checks are not billed per use, unlike some cloud provider billing practices, which can be very costly during a DNS DDoS attack.

Cloud scalability with F5

Operations like cloud bursting can enable an enterprise to scale its application resources in an on-demand basis and pay temporarily for the resources utilized during these occasional spikes. Effectively bursting to the cloud requires BIG-IP local traffic management, global server load balancing services, and BIG-IQ cloud orchestration that provide a number of very tightly choreographed operations occurring automatically, which are transparent to end users:

- When a preset threshold is reached, resources will be provisioned and deployed in the cloud.
- The new application server resources are automatically added to the pool of resources.
- Users are redirected to the application resources in the cloud.
- When application loads fall below the preset threshold, the cloud infrastructure resources are de-provisioned and all connections are redirected back to the data center.

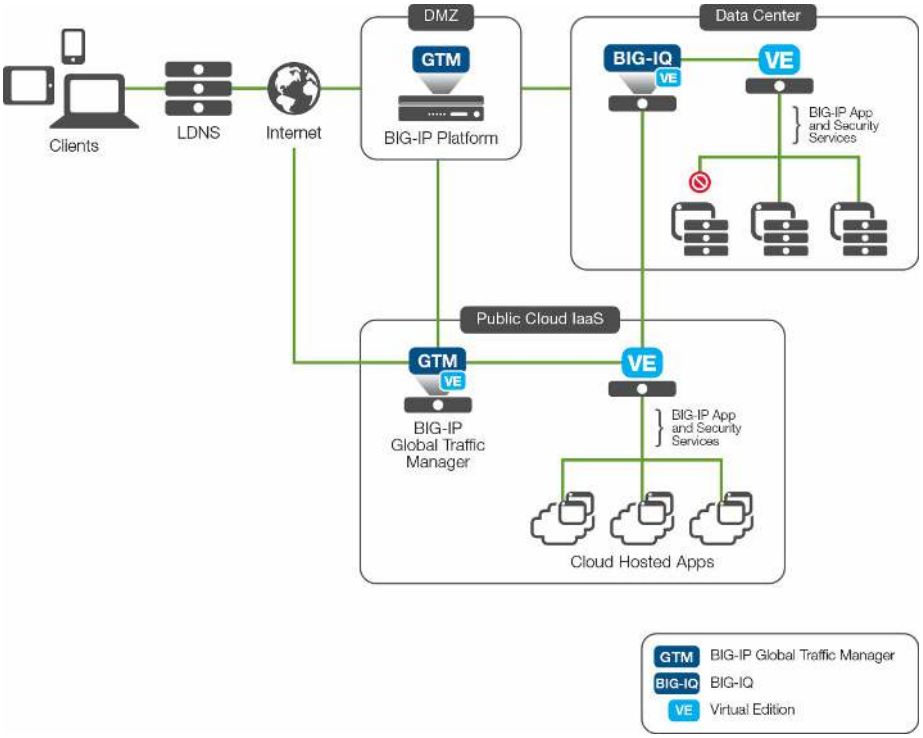

Cloud Bursting with F5 application services

Visibility, control, and consistent policies

F5 iApps® templates enable enterprises to deploy the application delivery services necessary to support their applications in minutes. iApps templates define the configuration and policies of services such as traffic management, encryption, firewall, and performance optimization for each application. iApps and inline analytics provide complete control and visibility into your applications for consistency and better troubleshooting.

In addition, F5 iRules® scripting language—F5's traffic scripting interface—enables programmatic analysis, manipulation, and detection of application traffic to enable customization and rapid response to zero-day threats, network conditions, and business requirements. The portability of iApps and iRules enables application mobility across and between cloud providers.

Management, Automation, and Orchestration

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 avoiding management silos.

The F5 solution provides integrated and automated application delivery capabilities into the cloud provider, rapidly reducing the provisioning and deployment times for application networking services. It accomplishes this through:

- Enterprise data center integration into third-party cloud management tools.
- Automation of the provisioning of application delivery services.
- Orchestration that expedites deployment times.
- Extensibility and unparalleled flexibility via the REST API.

BIG-IQ® is F5's intelligent management and orchestration platform and provides an open, programmable framework to manage physical and virtual BIG-IP solutions across private and public clouds. BIG-IQ helps organizations deploy and manage application and security services in a fast, consistent, and repeatable manner—regardless of the underlying infrastructure. In addition, BIG-IQ integrates or interfaces with the major cloud orchestration engines and cloud providers through REST APIs and cloud connectors.

BIG-IQ provides lifecycle management of iApps, which simplifies and automates the provisioning of application delivery services rapidly. The IT organization can define a catalog of available application delivery services, including customized or multi-tiered offerings, from which administrators and application managers can quickly select as needed.

BIG-IQ centralized management and orchestration

Flexible licensing models

BIG-IP virtual editions are available in Good-Better-Best bundles in utility billing (hourly, daily, monthly), annual subscription, bring-your-own-license (BYOL), and Volume Licensing Subscription models. For test/development pre-production workloads or temporary cloud bursting and scalability use cases, F5 offers 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. BYOL is ideal for hybrid cloud environments where you want to take existing BIG-IP VE licenses from your private data centers directly into the public cloud. The VLS offering is designed for enterprises 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.

	Non-production and Bursting	Production	Scale Production	
Business Scenarios	• Development/Testing/PoC • Temporary capacity needs	• Run-rate workload production • Disaster recovery • Longer term stable environments	• Large scale workload production • Hybrid cloud environments	
License	Utility (Hourly, Daily, Monthly)	Annual Subscription	BYOL (bring your own license)	VLS (volume licensing subscription)
Business Scenarios	Full architectural flexibility to deploy what you need when you need it and pay only for what was used	OPEX-based production with predictable cost	Optimized cost in the cloud while having flexibility to move licenses between clouds and DCs	Value cost structure that enables you to bring BIG-IP app service to all applications in the cloud and your DCs

F5 flexible cloud licensing options

Conclusion

Adoption of public cloud services has grown exponentially over the past few years. Many IT startups and some mature media companies have even deployed entirely in the public IaaS cloud providers with significant success. As enterprises undertake plans to migrate more critical applications to the cloud, the proven benefits of application delivery using the BIG-IP platform can easily be ported to cloud application workloads. Doing so will address many of the challenges and concerns that enterprise customers have regarding adoption of the public cloud including consistent security across all application infrastructures.

The F5 cloud migration solution leverages F5's BIG-IP application services and BIG-IQ products to deliver critical application availability, performance, and security—regardless of location. With flexible licensing models available in the leading cloud marketplaces, enterprises can plan, stage, and deploy these services to the public cloud. F5 enables enterprises and organizations to confidently transition workloads to single- and multi-cloud environments while maintaining visibility, security, and control.

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