



White Paper

Automating F5 Application Services for VMware with F5 BIG-IQ Cloud

Virtual machine and application provisioning are now being automated while network service provisioning is not, thus causing delays. Integration between the F5 intelligent services platform and VMware addresses this gap to allow policy-based provisioning of application networking.

By Simon Hamilton-Wilkes

Principal Solution Engineer



Contents

Introduction	3
<hr/>	
The App Delivery Provisioning Process	3
The BIG-IQ Cloud Platform	6
BIG-IQ Cloud and VMware	7
<hr/>	
Conclusion	9



Introduction

Expectations of enterprise IT have changed, and virtualization has already made silos of physical infrastructure hosting a single application obsolete. Self-service portals, with near instantaneous provisioning and lower costs, are the epitome of the new goal: cloud computing.

VMware is the leader in the private cloud space with its vCloud Director, now part of the vCloud Suite. This product builds upon VMware's vSphere virtualization platform to deliver the user portal, automation, and management tools required to make an Infrastructure as a Service (IaaS) cloud.

Application provisioning adds the next layer of automation to permit multi-tier applications to be deployed with a few clicks of a mouse. This automation brings a new challenge, however. Though virtual servers and applications may be created rapidly, the network and security services required to meet service level agreements (SLAs) and optimum use of the application may take far longer to deploy, as other teams may be responsible for security and application delivery services.

Manageability is also a key factor in efficient provisioning of virtual resources, since Gartner research indicates that about 80 percent of total IT spending in 2012 was for management. To reduce costs, it's highly desirable to also streamline management by reducing the proliferation of consoles and points of configuration. Improved manageability and automation also reduce costs by reducing human error at every stage.

The App Delivery Provisioning Process

Application delivery is the glue between users and their applications in the data center, providing opportunities to manage performance and security in addition to the availability and scaling benefits of traditional load balancing. Local traffic management is required to enable multiple servers to serve each tier of an application in a manner that ensures availability and the ability to scale out. Global traffic management extends these capabilities, taking advantage of DNS services to allow applications to span multiple physical or cloud data centers in an active/active or active/passive mode in order to dynamically direct users to the closest, least loaded, or most cost-effective data center according to policy.



For organizations taking advantage of F5 Application Delivery Optimization (ADO) solutions, security and acceleration are both derived from the application-layer proxy within F5® BIG-IP® products. This proxy permits access control and protocol validation in addition to TCP optimization, caching, compression, and SSL offloading, leading to a better experience for the application user and a lower load on the virtual servers. This proxy also permits the network to react dynamically and appropriately, depending on whether users are accessing locally from a high bandwidth connection or remotely via a phone or tablet.

Achieving those application delivery objectives depends on the correct configuration of all parts of the application delivery system, whether in the data center or in the cloud. One of the remaining barriers to agile utility computing in private clouds, however, is the separation of roles between the application owners and the compute, storage, networking, and security teams. While vCloud Director helps to consolidate management and partially overcome the inefficiencies and communication challenges of virtual resource provisioning involving multiple teams, the necessary links to incorporate application networking and security-related provisioning are the least developed. The existing provisioning process may be technically straightforward and yet still rendered complex due to “layer 8” organization and policy issues.

When the need for a new instance of an application is recognized and communicated to the IT organization, a multitude of processes are initiated. Self-service provisioning may be available for the virtual machines themselves. Some organizations also have template-based provisioning for operating systems and patches. Nonetheless, even in these situations, many different individuals or teams must pass off information and tasks to one another to allocate IP addresses and DNS, set up security policies, correctly configure application delivery, and then document and perhaps bill for those activities as appropriate. The handoff of critical information and the re-queuing of necessary tasks, particularly when most teams already have full workloads or significant backlogs, may extend the required provisioning time from hours to weeks.

Self-Provisioning via Template Services


Organizations relying on F5 ADO solutions, including BIG-IP products, can streamline this process, thanks to F5® iApps™ technology and integration between the F5® BIG-IP™ platform and vCloud Suite. The combined F5 solution enables organizations’ BIG-IP system administrators to publish template services to consumers of application services or to vCloud Suite administrators. This permits the providers (for example,



the security team) to retain control of the subset of features that can be self-provisioned by application owners, thus ensuring compliance with policy, while empowering application owners to select the features they need, since they are the ones most likely to know, at least in broad terms.

This template approach increases efficiency, speeds provisioning of application services, and reduces the risk of configuration errors by reducing the need for various specialists to perform repetitive, manual tasks for each new service instance. Common policies can be defined and standardized by each administrative owner and applied consistently. Because the specialists retain control of exactly what is presented in the templates, existing technical and organizational policies are not disrupted, eliminating a potential barrier to adoption. The publication of service templates also enables the creation and monetization of different classes of service or SLAs for charge-back or charge-through situations. Finally, templates offer the advantage of providing control over the visibility or opacity of many minor settings that may be undesirable to expose to application owners or service consumers because they are too complex, are prone to errors in selection, or simply add unnecessary complexity to the self-service process.

Without F5 products



With F5 products

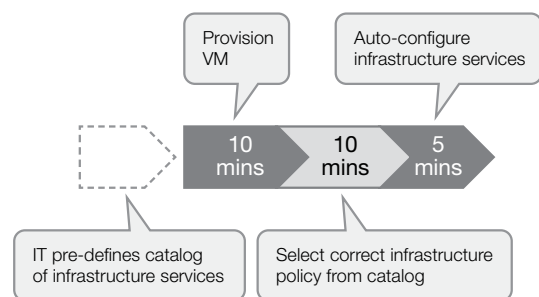


Figure 1: Application services provisioning with and without F5 solutions and vCEF



iApps Templates for Self-Provisioning ADN

The integration between VMware and F5 products relies heavily on iApps, an application-centric deployment packaging technology designed to encapsulate and automate the configuration of BIG-IP Application Delivery Networking (ADN) services for an application instance. A single iApps Template may set hundreds of parameters on a BIG-IP device or virtual edition, creating and configuring anything from basic virtual IPs and pool members to setting advanced parameters for an application deployment.

Examples of pre-defined iApps Templates that ship with BIG-IP® Local Traffic Manager™ (LTM) are those for Microsoft Exchange, Microsoft SharePoint, and VMware View. The result of extensive application understanding and testing, these iApps Templates are maintained as applications are incremented. iApps Templates can also be created and customized by customers or F5 Professional Services to suit bespoke applications or specific business needs.

In addition to greatly reducing configuration time, iApps Templates limit manual user input to eliminate human errors, increasing accuracy and consistency and thereby reducing the need for troubleshooting. Each iApps Template has input variable requirements, which can be fixed by the author or administrator or exposed as configurable parameters at deployment time, according to business policy or the degree of simplification desired. To deploy F5 ADN services for an application, the consumer simply selects the appropriate iApps Template from the available catalog provided by administrators, fills in the required parameters and customizations, and presses “Finish.” The iApps technology takes over from there, instantiating new application services and automating their full configuration, including creating necessary objects, setting parameters, and activating the service(s) specified. This streamlined workflow allows the configuration and deployment of a full ADN service profile, including the full spectrum of F5’s rich L4-L7 services, in a matter of minutes. The result is optimal performance, security, and availability for the application—fast.

The BIG-IQ Cloud Platform

iApps Templates are a key component of the new F5® BIG-IQ™ Cloud management platform, which enables fast, self-service provisioning of application services in environments featuring BIG-IP devices or virtual editions. Specifically, BIG-IQ Cloud provides a framework for orchestrating BIG-IP network services in public and private



clouds by facilitating the publishing of iApps Templates into the VMware user interface and by speeding subsequent instantiation and deployment.

The BIG-IQ Cloud management platform consists of:

- The BIG-IQ Cloud REST API.
- BIG-IQ Cloud Connectors.
- The BIG-IQ Cloud portal.
- iApps lifecycle management.
- BIG-IQ service health monitoring.

A BIG-IQ Cloud Connector plug-in bridges the BIG-IQ Cloud REST API to the VMware REST API to provide this integration.

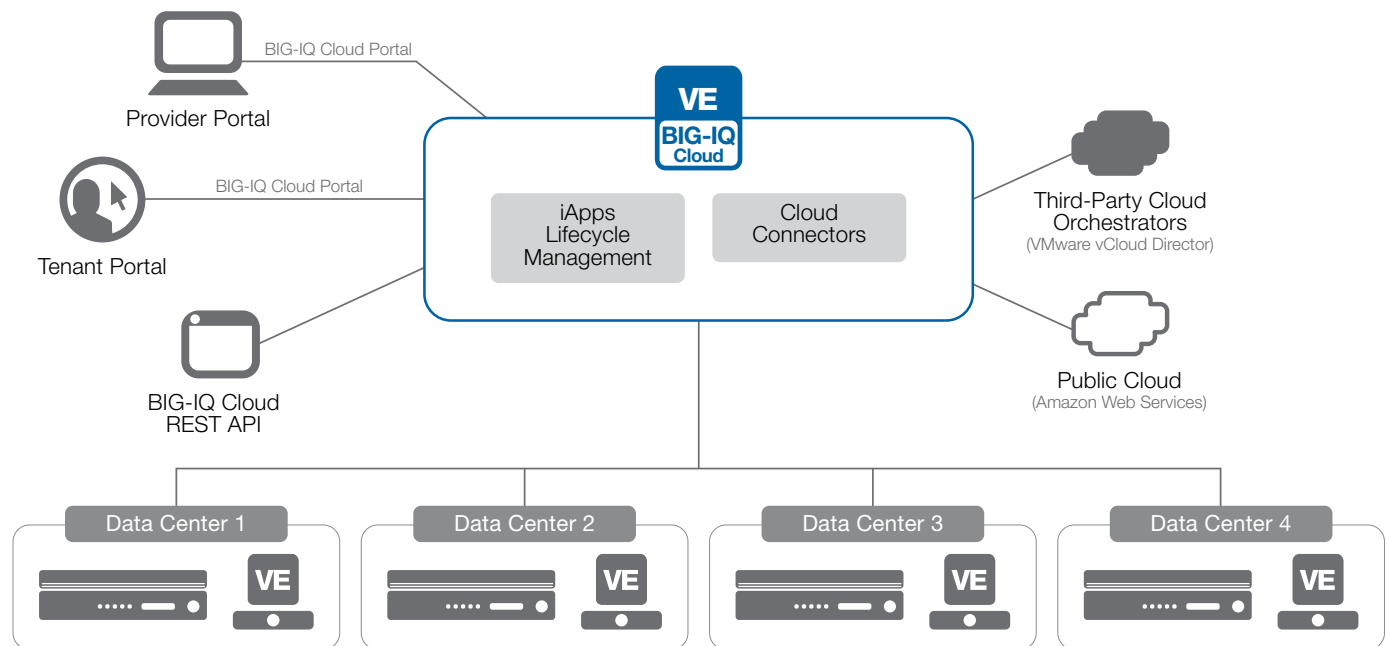


Figure 2: BIG-IQ Cloud APIs and multi-directional communication

BIG-IQ Cloud and VMware

Communication between VMware and F5 Application Delivery Controllers (ADCs) occurs via two-way, REST-based APIs. To expose BIG-IP ADC services within VMware's platform, BIG-IQ Cloud registers itself as a service manager and publishes available iApps Templates into a catalog natively accessible as an Edge Gateway



service from the VMware user interface. When an iApps Template is selected from that catalog and filled out to choose from the available options, the template and its parameters are pushed via REST to BIG-IQ Cloud, which in turn pushes them to all of the relevant BIG-IP ADCs and implements the proper configuration.

Since the iApps Template can configure any feature or function offered by the organization's BIG-IP platform, the complete F5 feature set is natively accessible for configuration and deployment through the VMware console. When a service is requested, automatic queries and responses between the different APIs replace the back and forth shuttling of tasks and information between different teams of people. This automation of otherwise repetitive, detail-oriented tasks enables greater scalability.

In addition to full configuration management, BIG-IQ Cloud provides centralized management services for BIG-IP device inventories, service contract renewal, and monitoring and proactive management of system health.

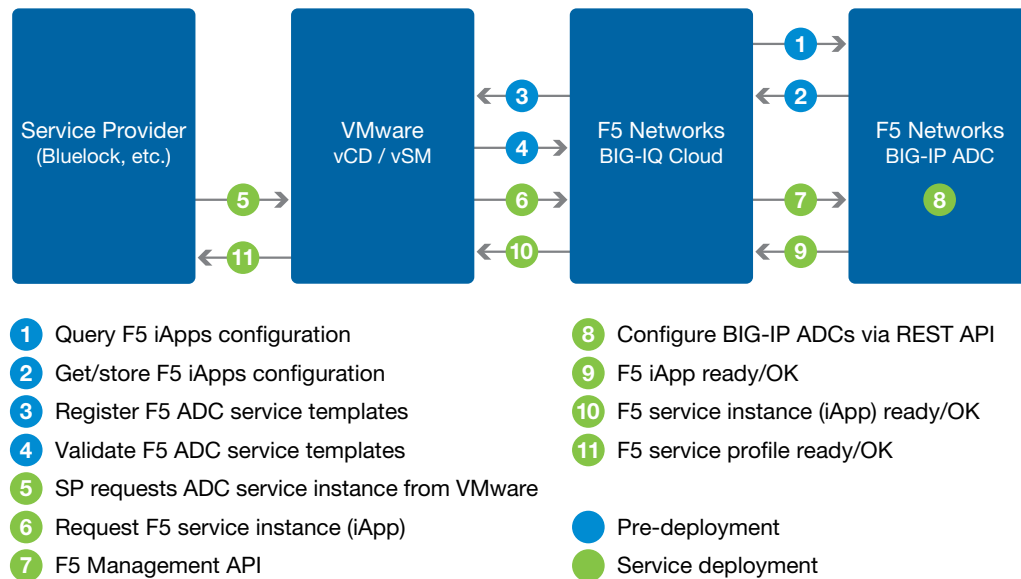


Figure 3: The F5 and VMware API-based service provisioning (or teardown) workflow

Conclusion

The integration of F5 BIG-IP ADCs with VMware products through BIG-IQ Cloud enables self-service provisioning of application networking in data center and cloud environments while retaining policy and configuration quality control for the teams responsible for security and application delivery services. Multiple policies and service tiers can be defined in the catalog of application service templates to provide both choice and consistency as well as greater automation and its benefits.

Automated provisioning of application services translates to increased virtual data center agility with fewer configuration errors. BIG-IQ Cloud removes repetitive manual tasks from the workloads of BIG-IP device administrators. It also empowers application owners to get the application networking services they need faster, with total configuration time reduced from days or weeks to minutes. The result is greater operational efficiency, fewer consoles, and lower management costs, freeing more time for skilled personnel to work on more challenging tasks.

