



vCenter Orchestrator Automates F5 BIG-IP Platform

Automate application delivery services, enable data center and cloud provisioning, and reduce deployment time with VMware vCenter Orchestrator's new plug-in for the F5 BIG-IP platform.

White Paper

by Simon Hamilton-Wilkes
Principal Solution Engineer
VMware Alliance



Contents

Introduction	3
<hr/>	
Automating Workflows with VMware vCenter Orchestrator	3
<hr/>	
F5 BIG-IP vCO Plug-In	4
<hr/>	
Using vCO with the BIG-IP Platform	6
Auto-scaling workflow	6
Web farm patching workflow	7
<hr/>	
Conclusion	8



Introduction

Expectations of IT have changed, and virtualization has 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. 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.

Application delivery services help organizations ensure that their applications are fast, secure, and available to users. Efficiently deploying applications in the cloud requires streamlined automated provisioning, which allows organizations to avoid falling back on time-consuming manual processes. Improved manageability and automation also reduce costs by reducing human error at every stage.

With VMware's vCenter Orchestrator (vCO), organizations can automate and optimize virtual infrastructure deployment in the cloud. And a new plug-in for the F5® BIG-IP® platform allows vCO to directly speak to a BIG-IP API, giving vCO the added capability of sending commands to and receiving information from BIG-IP devices.

Automating Workflows with VMware vCenter Orchestrator

VMware vCenter Orchestrator (vCO) links actions into repeatable sequences, called workflows, to automate tasks. vCO provides a drag-and-drop visual interface, and includes pre-built building blocks to perform common tasks, as well as a library of available plug-ins to integrate with other systems. This enables vCO to send emails, connect to databases, and use command-line tools to control other systems where APIs are not available.

In addition, vCO has a REST API itself—allowing it to be leveraged as an orchestration engine by other management tools, both from VMware and third parties. The F5 BIG-IP platform has its own automation API, iControl®. A new plug-in module allows vCO to speak to iControl, which adds the capability of sending commands to BIG-IP devices, and also interrogating them for statistics or configuration information.



As a full-proxy solution, the BIG-IP platform provides the glue between users and applications, with the ability to respond dynamically to the context of the session, who the user is, what device they're using, and where they're coming from—and tailor the connection parameters to suit. In addition to load balancing and DNS services, acceleration and security features can be provided within the appliance (virtual or physical)—offering a full suite of application delivery features.

The downside to all this capability is complexity, which is addressed by F5 iApps® templates, which pre-define the optimal settings for an application to allow fast, consistent, and repeatable deployments. The F5 BIG-IQ™ Cloud module takes this further by providing a mechanism to pre-publish specific iApps to selected users within a cloud infrastructure. Parameters can be marked as user editable or not, to accommodate security and business policy. Consolidating the application parameters simplifies the deployment, abstracting the details while retaining the breadth of BIG-IP features. BIG-IQ Cloud also supports API integrations—northbound via its own API and eastbound via cloud platform API connectors.

VMware vCloud Automation Center (vCAC) allows organizations to automate the provisioning and deprovisioning of servers, applications, and services—and gives a policy blueprint defining who is permitted to use what resources. The self-service portal can be tuned to be more or less complex to fit the target user base. vCAC integrates with vCO as it dynamically discovers workflows available to it via the vCO REST API. This multilevel approach to automation permits organizations to separate business policy definition and enforcement from the details of implementation. Complex applications that require many tiers of services to be instantiated and coordinated can be supported, with the BIG-IQ and BIG-IP platforms providing application delivery and security services.

F5 BIG-IP vCO Plug-In

The vCO plug-in supports three different connection methods to permit different BIG-IP software versions, using SOAP or REST APIs, and the use of the BIG-IQ platform. When connecting directly to a device, workflows can use the individual actions included with the plug-in, or build up compound tasks where many commands are passed by a single workflow.

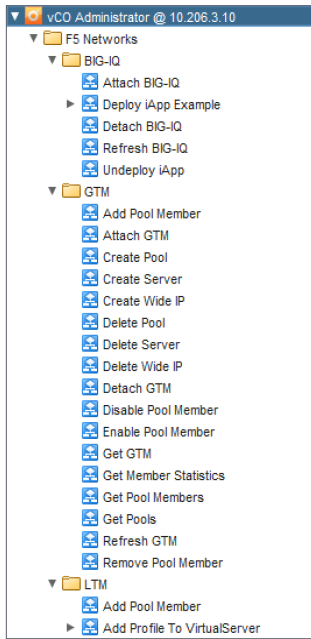


Figure 1: List of actions in plug-in package.

When connecting via the BIG-IQ platform, the plug-in first determines which iApps templates are available, and then presents these as choices, capturing the required variables. The cloud administrator will have already chosen which of these choices are editable and which are static, allowing organizations to simplify the deployment options and enforce security or business policies. Using the BIG-IQ platform as a point of policy control combined with vCAC cloud policy blueprints can ease organizational issues around self-service provisioning by application owners.

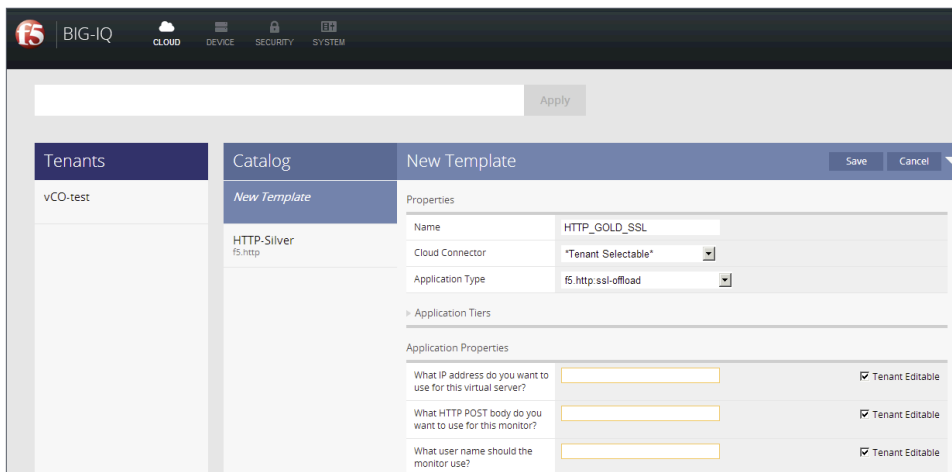


Figure 2: Application templates in BIG-IQ Cloud simplify service deployment.



Using vCO with the BIG-IP Platform

Provisioning application delivery services is the primary and obvious use case. In addition, the vCO plug-in can help organizations deprovision application delivery services, which enables life cycle management of load balancing, DNS services, and security configuration—removing redundant policies when applications are modified or retired. It can simplify system maintenance, as pool members are disabled and enabled automatically to prevent service interruptions while servers are patched, or redirect users through DNS changes prior to a planned data center outage.

More ambitiously, organizations can configure closed-loop remediation where a workflow is initiated by another tool—such as VMware vCenter Operations—as a result of an error being generated or a threshold crossed with correlation occurring both prior to vCO and during the workflow. Role-based access control combined with the vCO or vCAC portals provides a tool for allowing non-networking and security personnel to interact with organizational BIG-IP devices in a controlled and user-friendly fashion.

Auto-scaling workflow

This is an example of a schema, which is the graphical representation of a vCO workflow. Here, a workflow auto-scales a web farm and its associated pool on BIG-IP® Local Traffic Manager™ (LTM) according to the established connection count.

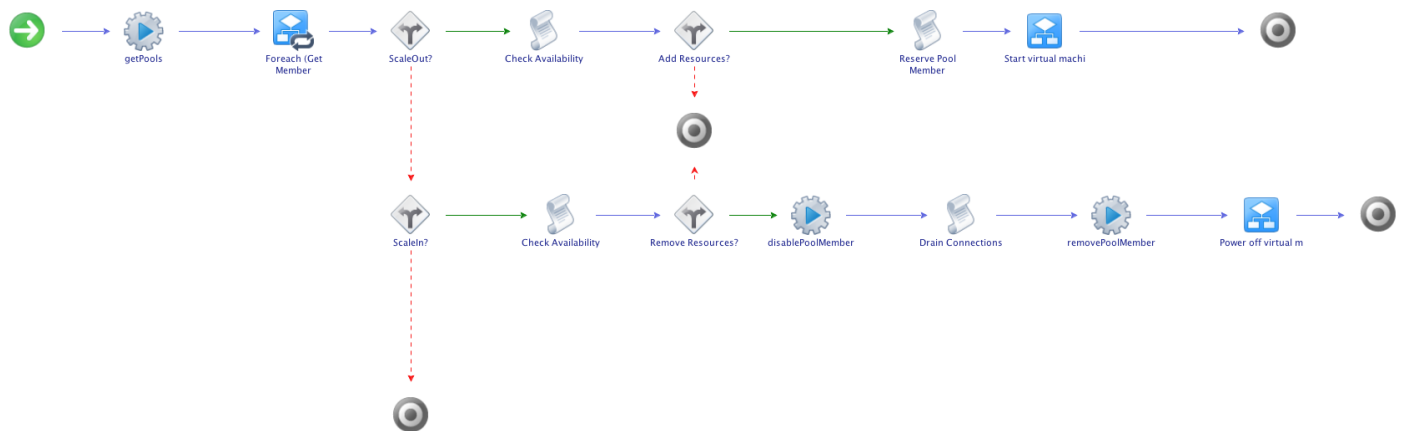


Figure 3: The schema for an auto-scaling workflow.

The workflow begins by fetching pool member statistics—then it can make a scale-out or scale-in decision based on the member connection counts.



If it decides to scale out, then it uses the vCenter plug-in to see if there is another web VM available to add; if not, the workflow ends. If there is, it sends a power-on instruction via the vCenter API and also an instruction to BIG-IP LTM to add the new pool member. The workflow waits for the power-on action to return from vCenter—which takes a moment as vCenter waits to ensure VMware tools are running, indicating the virtual machine is healthy—and then finishes the process.

If it decides to scale in, the workflow starts by disabling the highest-numbered pool member and waiting for connections to drain gracefully. When no connections are active to that member, it issues a standby instruction for that virtual machine via the vCenter API and removes the pool member from BIG-IP LTM, then finishes the process.

While in this case only F5 and vCenter API calls are used, it could easily be expanded to send an email to report the capacity changes it makes or to warn if it exhausts the available web server VMs. This workflow could also be set to run on a schedule, or be started by an external process such as vCenter Operations if an alert is triggered about web page latency, for example.

Web farm patching workflow

In this case a workflow could be initiated manually or by schedule in off-peak hours to install maintenance patches on a web cluster without human intervention or user impact.

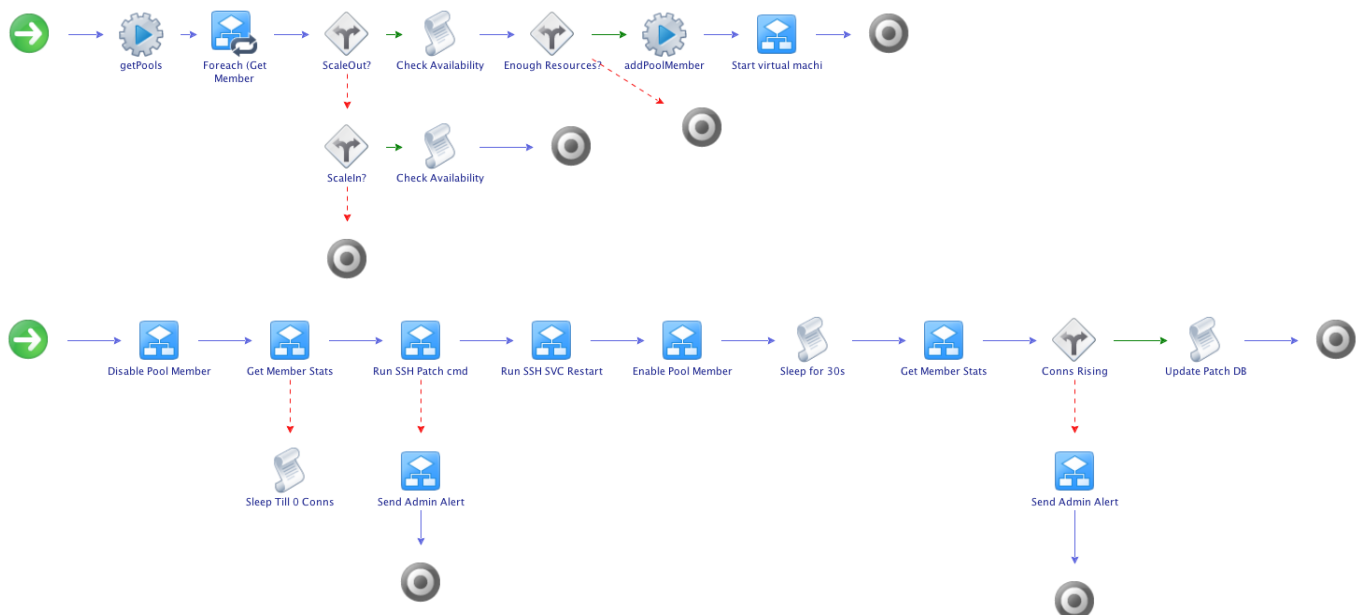


Figure 4: The schema for a maintenance workflow.

Initially the workflow checks that it has patches to apply, and that the connection load is below a pre-set threshold for off-peak. It then picks the lowest-numbered web server and disables it from the load balancing pool, to allow connections to gracefully drain. Once the connection count has dropped to zero, it connects to the server, installs the patches, and restarts services or reboots. It then enables the server in the pool and waits for health check to pass and the connection count to be rising before looping to repeat on the next server.

Conclusion

Orchestration of application services translates to increased virtual data center agility with fewer configuration errors. The result is greater operational efficiency, fewer consoles, and lower management costs, freeing more time for skilled personnel to work on more challenging tasks.

The F5 vCO plug-in enables the automation of BIG-IP services from a VMware environment, both manually from a self-service portal or web client, and automatically to enable a self-healing network.

This makes end-to-end data center and cloud provisioning possible—including application delivery and implementation of security policies—reducing the time to deploy a new application environment to minutes rather than days or weeks. Other use cases such as maintenance can be automated too, for a reduction in human intervention, which leads to greater repeatability and less chance for error.

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

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com



Solutions for an application world.