



F5 iControl

In this white paper, get an introduction to F5 iControl service-enabled management API.

White Paper
by Lori MacVittie



Introduction

As applications have become increasingly dependent on the network and its infrastructure for scalability, reliability, and performance it has become necessary to ensure that the infrastructure is capable of reacting dynamically to the ever-changing conditions inside the data center that can adversely affect the delivery and performance of applications.

One of the core components of such a dynamic network infrastructure is an Application Delivery Controller (ADC). ADCs evolved from load balancers and now include a plethora of functionality and features that enhance the performance of applications while providing unparalleled availability and reliability. However, simply being capable of adjusting to changing conditions and recognizing those conditions is not always enough. Sometimes it is necessary to go further, to integrate the application and the network infrastructure more tightly in order to provide greater value for application stakeholders.

With the advent of cloud computing concepts, whether deployed "in the cloud" or "in the data center," it has become even more important to tightly integrate the network infrastructure into the workflow processes that enable, disable, create, remove and migrate applications in a variety of containers both virtual and non-virtual. For F5's application delivery platform, that integration is realized through the use of iControl.

iControl Basics

What is iControl?

iControl is a Web services-enabled open API providing granular control over the configuration and management of F5's application delivery platform, BIG-IP. iControl, like other SOA and Web services-enabled solutions, can be used by virtually any platform capable of integrating via SOAP (Simple Object Access Protocol). It can be used to build custom management and monitoring applications, to integrate with business process management (BPM) and other workflow applications, and can be integrated directly into applications to provide better control over the delivery of the application. It also integrates with virtual computing platform management tools like VMware vCenter and Microsoft System Center Virtual Machine Manager (SCVMM) to help orchestrate the automated provisioning and de-provisioning of applications that can be used to build private and public cloud computing environments. The iControl interfaces enable access, based on management access policies, to all configuration and management policies of F5's application delivery platform.



WHITE PAPER

F5 iControl

Using iControl you can control the way your application is delivered based on factors you define. You can modify the F5 BIG-IP device's configuration in myriad ways, including:

- Adding servers dynamically to an application pool
- Stopping requests from going to an application server
- Changing the way requests are routed to servers
- Influencing the choice of servers based on current application or server load

iControl applications have been implemented using a variety of languages and development environments for use on a number of different operating systems and platforms. Because iControl is standards-based, it is interoperable and usable by any language and on any system capable of communicating via SOAP Web services.

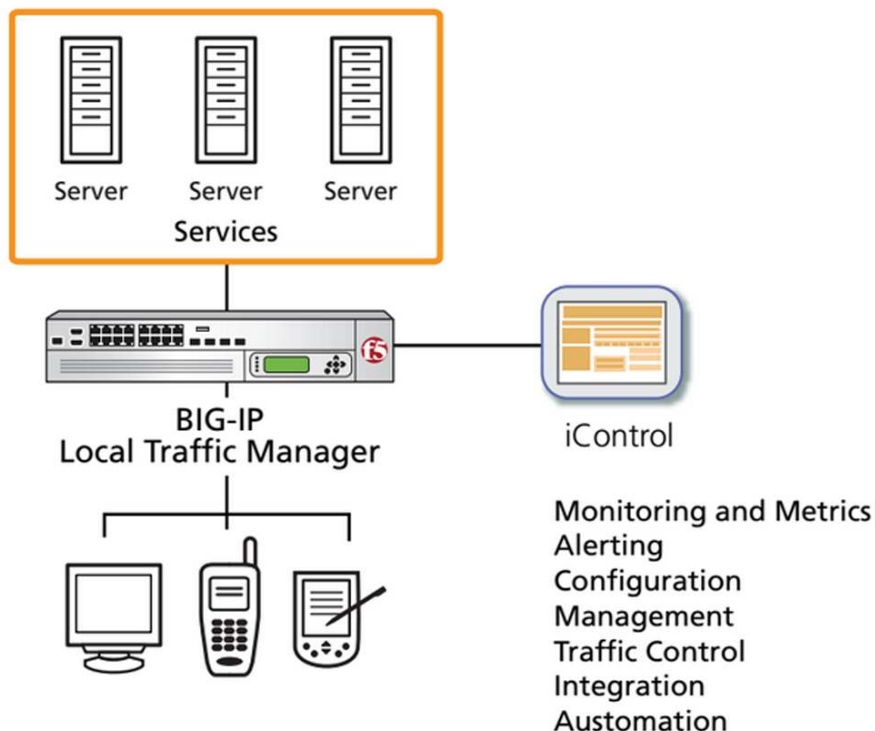


Figure 1: iControl services

Assemblies and libraries that make it easier to use iControl with specific languages and environments have been developed and are available for download from DevCentral, F5's community of developers, architects, and administrators. Java and .NET are the most widely supported, but PERL, PHP, PowerShell, and Python all have examples or libraries to aid in development of iControl applications.



WHITE PAPER

F5 iControl

Architecture

The iControl interfaces are accessible via standards-based WSDL (Web Services Description Language) version 1.1. using the F5 BIG-IP management interface. iControl presents the BIG-IP system as a set of services, each providing a broad range of configuration and management capabilities.

The WSDL interfaces can be accessed via HTTPS and can then be used as any other service description within applications or third-party products such as BPM systems, ESB (enterprise service bus) or SOA governance registries.

The options for managing and monitoring BIG-IP devices are vast, and therefore there are multiple WSDL documents, each targeting specific namespaces (a way of uniquely identifying objects so that they do not clash when used together) under which a grouping of like operations are contained, such as: management, networking, monitoring, and local load balancing.

This enables you to integrate only the interfaces you require, much in the same way you import only specific packages or libraries during development. The majority of the iControl interfaces are designed for F5's core application delivery platform, the BIG-IP system, though additional service interfaces are available for add-on modules and products designed to be deployed on BIG-IP system, such as secure remote access and web acceleration.

iControl interfaces are accessible only via an SSL-protected connection, and require HTTP basic authentication to prevent unauthorized access.

iControl Use Cases

Because iControl is a Web services-enabled API, there is virtually no limitation on how it can be used to interact with F5's application delivery platform. If you can think of a solution then it is likely you can build it. But in order to understand how iControl might be useful in your environment, here are some of the ways in which customers and partners have used iControl to add value to their architectures and applications.

Data Center Automation

Cloud computing providers taking advantage of virtualization aren't the only companies automating their data centers. A wide variety of organizations that are taking advantage of the benefits of operating system and server virtualization are looking for ways to better manage their shiny new virtual infrastructure and are finding a lack of tools capable of doing so.



WHITE PAPER

F5 iControl

While virtualization environments are generally service-enabled, making the process of deploying and moving virtual instances of applications a breeze, the rest of the supporting infrastructure is often not service-enabled, requiring intense manual processes in order to manage a new virtual infrastructure.

iControl solves this problem by integrating into third-party applications that manage virtual environments. The most common scenario is on-demand application scalability achieved through the automated launching of new application instances when demand for that application increases. Using iControl, F5's application delivery platform can be instantly notified of the new instance of the application running on a new virtual machine and begin directing requests to that instance. Similarly, when demand lessens and the additional application instances are no longer needed, the F5 application delivery platform can be notified through iControl to quiesce the current connections so the instance can be deactivated and the resources reapplied to other applications.

iControl can even take this automation further by monitoring application demand and recognizing when an application is about to be overwhelmed. When the monitoring application is notified that demand will overload an application, it can automatically initiate the workflow that will bring another virtual instance of the application online and notify the F5 application delivery platform. It can continue to monitor demand until it senses demand has decreased and initiate the quiescence, eventually deactivating the virtual instances on its own.

While iControl can be used to automate virtual instances of applications, it can also be integrated into applications to provide the same kind of automatic management that is offered by third-party solutions controlling virtualization servers. Developers can essentially add a "hook" to their application initialization routines that notifies the F5 application delivery platform that a new instance of the application is available and it should be added to the appropriate virtual server. For example, the Java Servlet specification 2.3 enables applications to define a listener specifically for servlet context. This listener generally responds to two events: initialization of the application and destruction of the application. By implementing a listener for the application, the F5 application delivery platform can be automatically notified of instances being created and deleted and add them to the appropriate virtual server. Similarly, when the application instance is destroyed, a message can be sent to the application delivery platform to remove the application instance from the pool.

Using iControl in this scenario offers great improvements in productivity, as what is often a manual process can be easily automated.



Proactive Exception Handling

F5's application delivery platform excels at intelligently monitoring the health of applications and adjusting the flow of requests based on the responsiveness and health of each server. While BIG-IP devices can detect and react to application errors and exceptions, it is often the case that it would be preferable to react before a situation makes a server unavailable or begins throwing exceptions.

One way in which this might be achieved is by communicating with F5's application delivery platform when the first signs of trouble are detected rather than waiting until the server or application is in full-blown failure mode. For example, assume an exception is thrown in the application that indicates that a resource of some kind—database connections, memory, disk—is running low, or that there is an impending problem. While it is quite possible that resources will be recovered or the problem resolved, it would likely be better for performance and availability if requests were redirected to other servers until the application recovers, or at least reduced in volume for a time. When the exception is caught, the developer can notify the BIG-IP device to reduce the priority of the server such that fewer requests are sent to that instance of the application, or stopped completely. This gives the application time to recover or, if necessary, for the administrator to take corrective action without concern that such action might interrupt service to that application.

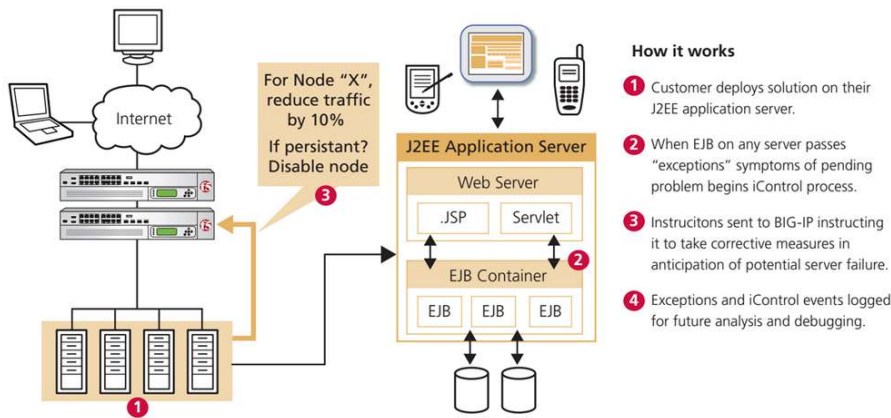


Figure 2: Proactive exception handling solution using iControl and the BIG-IP system

Conclusion

In a dynamic and often virtualized environment the ability of the application delivery infrastructure to adapt and integrate applications their management tools is a requirement for success. It is no longer possible in a dynamic environment, with its high rate of change in both network and application layers, for the network infrastructure to remain static.

WHITE PAPER

F5 iControl



iControl provides the means by which F5's BIG-IP platform can be automated, integrated, and flexible such that it can adapt in real-time to changing network and application infrastructure conditions inside and outside the data center. Whether you are building out a new infrastructure or adapting to new technology, using iControl can enable unprecedented levels of control over and visibility into the application delivery process.

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