



# The F5 Application Services Reference Architecture

Build elastic, flexible application delivery fabrics that are ready to meet the challenges of optimizing and securing applications in a constantly evolving environment.



## WHITE PAPER

The F5 Application Services Reference Architecture

## Introduction

In today's fast-moving, on-demand culture, users expect their applications to always work—and work fast—on any device. If organizations fail to meet these expectations, it has never been easier for customers to find alternatives. For IT departments or organizations to stay connected to their customers, they must meet user expectations with responsive, omni-channel applications that deliver new features seamlessly.

In addition to traditional, linear application delivery from corporate data centers to PC-based users, organizations can now deliver applications from a complex matrix of locations—including SaaS providers—outside the data center. Increasingly mobile users around the globe now demand this anywhere, any-device access.

At the same time, the pace of global innovation is always accelerating with faster release cycles, the emergence of DevOps, and the move toward continuous delivery. Although this paradigm shift better aligns IT and business, it also creates challenges for traditional infrastructures. The Application Delivery Network must promote innovation—but not at the expense of stability, security, or performance.

With F5, organizations can build elastic, flexible application delivery fabrics that are ready to meet the challenges of optimizing and securing applications in a constantly evolving environment.



## Business Challenges

If the expectations of users are simple and the need to innovate is clear, the solutions to these challenges are neither. Even today's best-run organizations are challenged by the rate at which IT is evolving, the growth of applications, and the complexities created by mobility, cloud computing, and virtualization.

To be successful in delivering applications, IT organizations must:

- Optimize application performance from anywhere and to anyone and any device—without increasing cost or complexity.
- Enable rapid deployment and innovation.
- Flex and scale on demand.

How can IT fulfill these requirements for innovation and adaptability while maintaining the reliability users demand? How can you drive revenue, service customers, comply with regulations, and reduce costs if you can't scale application services when you need to? The inflexibility of traditional infrastructures is a significant obstacle on the path to an agile, application-centric network. A programmable, scalable application delivery fabric can provide a layer of application services to help bridge this gap.

## Architecture Solution

To succeed in this dynamic environment, organizations must create networks that combine core layer 4–7 functionality, programmability, and application fluency into a scalable Application Delivery Network.

While these components all bring their individual benefits, the real, lasting value lies in the resulting synergy of the combined solution. Failure to deliver even one of these technology components greatly reduces the business value of the infrastructure.

## Core Functions

One of the main benefits of adding application services to your network is the rich set of functions you can leverage to improve the user experience, add security, or provide analytic data when application traffic passes through your infrastructure. While the core features of high availability and load balancing are required for most applications, there is a large range of additional benefits that the application delivery network should provide.

Application Delivery Controllers (ADCs) offer hundreds of functions, but they can be broken down into broad categories:

### "Good Enough" and "Lots of Dumb"

"Good enough" application delivery solutions are typified by Application Delivery Controllers offering high-traffic capacity and some layer 7 functions at a low cost point. Lacking application fluency, programmability, and true scalability, these low-end solutions are likely to constrain innovation, rather than promote it.

The "lots of dumb" paradigm of creating tens, hundreds, or thousands of virtual load balancing instances generally solves problems of scale and deployment speed, but lacks the functions or fluency of a true application delivery fabric.



## WHITE PAPER

### The F5 Application Services Reference Architecture

Protocol	Parsing, translating, and manipulating network and application protocols. Examples: IPV4 to IPV6 translation, HTTP to SPDY gateway services, FIX traffic parsing, SIP Call ID extraction.
Encryption	Encrypting and decrypting application traffic for security, inspection, and server offload. Examples: SSL decryption, SSL re-encryption, HTTP cookie encryption, SSL traffic visibility, and key protection.
Context	Providing contextual information about the application traffic. Examples: Device identification, location awareness, reputation awareness, connection metrics, and application and server health monitoring.
Optimization	Altering some aspect of the application traffic or connection to improve the performance of the application. Examples: HTTP object compression, TCP connection multiplexing, adaptive traffic distribution based on server health scores, and TCP protocol optimization.
Traffic	Controlling and manipulating traffic. Examples: Load balancing, high availability, connection limiting, content- based routing, and rate shaping.
Content	Parsing and manipulating application content. Examples: Stripping credit card numbers from server responses, and reading and logging FIX transaction IDs.
Analytics	Providing statistics and logging for application traffic. Examples: HTTP application page load times, logging application events, and client-side latency.

However, a huge array of capabilities alone is not enough. With hundreds of features that can be deployed in thousands of combinations, how can an organization best deliver an application? Testing every feature combination and analyzing the results is impractical. The technology must be equipped with built-in application knowledge. In other words, application fluency is mandatory.

## Application Fluency

In this complex application delivery environment, you must be able to customize delivery policies based on application and device. A lean, agile policy designed for capacity and availability management of a high-frequency financial trading application should be different from a more extensive, feature-rich policy for an application such as Microsoft Exchange. Attempting to use the same policy for these very different applications would result in neither application getting the services it needs.



## WHITE PAPER

### The F5 Application Services Reference Architecture

The two different applications each demand a specific combination of the many functions available from the ADC.

However, the labor-intensive process of handcrafting and testing customized application delivery policies can cause operational costs to skyrocket. Customizing delivery policies without adding complexity or administrative overhead requires a solution that has a broad set of functions, as well as the knowledge to deploy them in the right configuration. Repeatable, reliable templates for application delivery should be used to easily create application-specific services. Using pre-defined templates removes the overhead of managing complex configurations and generates tremendous payback in terms of saved time and increased reliability in application deployments.

## Programmability

In an environment where the only constant is change, the need to create new solutions and then rapidly prototype, test, and deploy them demands an infrastructure that is innovation-ready. Programmability of both the data plane (application traffic) and the control plane (device configuration) allows you to create adaptable solutions that meet your needs without waiting for new features to be developed by manufacturers—or to create novel methods of traffic distribution for a new application. Full programmability makes it simple for organizations to create software defined application services on top of any network infrastructure.

Integrating with monitoring and orchestration systems enables task automation, such as managing traffic for new application release cycles, deploying new services, or responding to changes elsewhere in the infrastructure. It removes delays caused by manual processes and provides an agile delivery environment where limited control of application traffic can be delegated to application deployment teams.

## Scalability

The fabric delivering the application services needs to be as rich, flexible, and extendable as the services themselves. This means scaling not just locally in terms of capacity, but also in terms of location and architecture—both across traditional infrastructures and into public and private clouds. To achieve true scalability, platforms must be secure enough to be connected pervasively throughout the network or into the cloud.

## The F5 Architecture Solution

The F5 architecture solution for application services allows organizations to create a robust and scalable application delivery fabric. This architecture revolves around some simple principles:

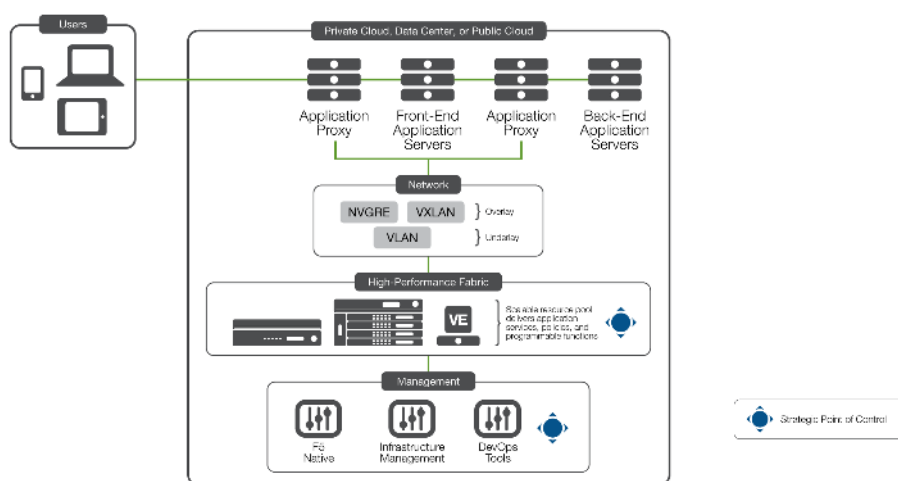


## WHITE PAPER

### The F5 Application Services Reference Architecture

- Create a services fabric of scalable and interconnected devices.
- Pervasively connect the fabric to the application and client networks.
- Use repeatable templates to define policies that deliver application-fluent services to optimize applications via application proxy points.
- Open up management and control of the fabric to orchestration and integration tools.

This creates a fabric that is able to deliver the right application services to the right applications at the right point in the network, for any application and any network.



The F5 Application Services architecture

At the heart of this fabric is the F5 BIG-IP platform, which is designed to supply application security, acceleration, and availability services. Available in physical, virtual, and cloud editions, the BIG-IP platform is easy to deploy into as many locations and architectures as required, while maintaining consistency across environments.

## Comprehensive L3-7 Functionality

From fine-grained manipulation of TCP protocol settings through more than 10 high-level application-specific proxies, the F5 BIG-IP platform provides a massive range of features and functions designed to optimize and deliver applications to users. The range of functions across all categories gives you the ability to create and use application services that are truly optimized for your applications, networks, and clients.



## WHITE PAPER

### The F5 Application Services Reference Architecture

## Application Fluency

Creating optimized delivery policies tailored for particular applications requires significant time and resources. With the BIG-IP platform, the work to build an application-fluent solution has already been done. The result is a suite of implementation services, templates, and guides designed to build best-practice configurations out of the box.

Application-specific proxies provide parsing and manipulation of protocols such as HTTP, SIP, and FIX. The BIG-IP platform's deep understanding of application protocols enables better traffic management and reporting—and opens up opportunities to control and manage traffic more effectively. Being able to route, drop, or manipulate traffic based on application values such as HTTP cookies, FIX message type, or SIP header makes it easier to simplify application infrastructures and build additional application capability into your delivery fabric. Having understanding and control of application traffic at a strategic point in the network can speed and simplify application development cycles. You can outsource application functions to the network, where they can be done more efficiently at a strategic point of control within the delivery stack.

F5 iApps Templates significantly cut deployment time and remove much of the potential for human error by providing a simple interactive process to configure and update application deployments. When you use an iApps Template, selection and configuration of all the required application functions is collapsed into a single, simple application-focused interface. Once an application has been deployed, the iApps system provides an application-focused view of the health and performance of the components that make up the deployment, which makes operational management easier and can provide feedback into the development or revision process for the next release. F5 maintains a library of new and updated iApps Templates and maintains partnerships with leading software vendors to ensure that templates closely follow new releases. No need to wait for software updates to manage new applications; simply download an updated template. iApps Templates are also fully customizable and designed to work in different deployment scenarios since your application delivery network should enhance your application, not constrain it.

Where iApps Templates are not available or suitable, F5 offers more than 100 application or solution deployment guides, which present detailed configuration information and best practices.

With a platform that understands applications and comes equipped with the knowledge to deliver them, organizations can save significant time and cost in deploying new applications and maximizing the value from existing ones.



## WHITE PAPER

The F5 Application Services Reference Architecture

### A Programmable Platform for Innovation

The F5 BIG-IP platform has a well-documented suite of programmable components across the control, data, and management planes, with which organizations can build an adaptive and automated application delivery fabric with real-time control over application traffic and device configuration.

#### Data Plane Programming

The F5 iRules scripting language offers real-time control of application traffic as it passes through your infrastructure. With the ability to route, reject, modify, and inspect traffic in any direction, organizations can build or deploy solutions to meet just about any application traffic challenge, no matter how unusual. iRules can use contextual application services such as IP reputation or location to make traffic management decisions—or manipulate server responses to remove sensitive information.

#### Control Plane Programming

The F5 iCall scripting framework allows the configuration of the BIG-IP platform to adjust based on data plane events, such as DHCP requests from new server instances or altering traffic distribution based on a change in monitoring data.

#### Management Plane Programming

The F5 iControl comprehensive management API enables integration with data center management frameworks, orchestration tools, and third-party applications and scripts. iControl also supports an event-based model, allowing applications and frameworks to subscribe to BIG-IP system events such as the change in status from up to down for a given application node. These three planes of programmability increase agility, reduce operational costs, and improve application deployments, freeing your network to promote innovation instead of constraining it.

### F5 and DevOps

To better align IT to business needs, organizations are moving to compressed application revision schedules and "continuous delivery" methods. Innovative DevOps practices automate the testing and deployment processes, driving consistency and faster times to production. The BIG-IP platform provides the application traffic control and the programmability required to realize these benefits within the production delivery network.





## WHITE PAPER

The F5 Application Services Reference Architecture

### A Secure, Elastic Delivery Fabric

Programmability, fluency, and functionality combine to create application-specific services that improve end user experience, increase security, and drive innovation. This is true only if they can be delivered at the right point in the network and at the scale required by the application. F5 provides a secure, elastic fabric that can deliver application services anywhere they are required.

The secure BIG-IP platform enables devices (physical or virtual) to be connected to the core network in a data center, or in a public cloud, to provide application services and managed control points across all zones. Functionality and performance are no longer locked in silos of devices in particular zones, but are available at all points in your infrastructure. This improves device utilization and reduces implementation costs for new applications, as the services can be supplied from existing resources no matter where the services are required.

The F5 ScaleN architecture creates clusters of up to 32 devices (physical, virtual, or mixed), which form a highly available, elastic delivery fabric. Application workloads migrate between devices in the event of failure and can be moved should more resources be required. Devices can be upgraded physically or by license key to process more traffic or enable additional features. The ScaleN fabric allows the incremental addition of capacity and services, enabling a smooth growth and cost curve as traffic increases. Workloads can be managed to run on the most appropriate device, and additional resources can be provisioned quickly via virtual editions, giving organizations fine-grained control of resources and putting the administrator in control of applications.

Network virtualization support in the BIG-IP platform allows the unification of virtual infrastructures and high-performance network hardware. Support for VXLAN and NVGRE (on virtual platforms only) enables the delivery fabric to extend from physical devices into virtual networks, bringing the power of specialized hardware for compression, encryption, and network processing into the virtual data center.

### Business Benefits

By creating a flexible, scalable, application-fluent network, organizations can optimize applications to a diverse user population and develop and deploy applications faster. To the IT organization, this means greater user satisfaction, better alignment with the business (through more agile application releases), and the ability to scale infrastructure in greater alignment with needs.

What are the benefits to the enterprise as a whole?



## WHITE PAPER

### The F5 Application Services Reference Architecture

A better user experience ensures that visitors become customers and employees become more productive. Effective delivery to mobile platforms increases the reach and effectiveness of applications, delivering more return on the acquisition or development investment.

Increasing the pace of application deployment cycles ensures that projects can make smaller adjustments to functionality to react to the analytics of application use and functionality. Projects that might otherwise fail to deliver can be corrected faster and with less cost. Using standard functions from the application delivery fabric reduces coding and costs.

The scalable fabric allows CapEx to align with growth, improving cash flow and reducing risk. This efficient application delivery fabric also removes islands of capacity, because services are delivered from a pool of resources, reducing overall spend through better utilization.

## Conclusion

A scalable, programmable, application-fluent fabric allows organizations to meet user expectations, innovate, and scale in a cost-effective manner aligned with business goals.

F5 Networks, Inc.  
401 Elliott Avenue West, Seattle, WA 98119  
888-882-4447 [www.f5.com](http://www.f5.com)

Americas  
[info@f5.com](mailto:info@f5.com)

Asia-Pacific  
[apacinfo@f5.com](mailto:apacinfo@f5.com)

Europe/Middle-East/Africa  
[emeainfo@f5.com](mailto:emeainfo@f5.com)

Japan  
[f5j-info@f5.com](mailto:f5j-info@f5.com)