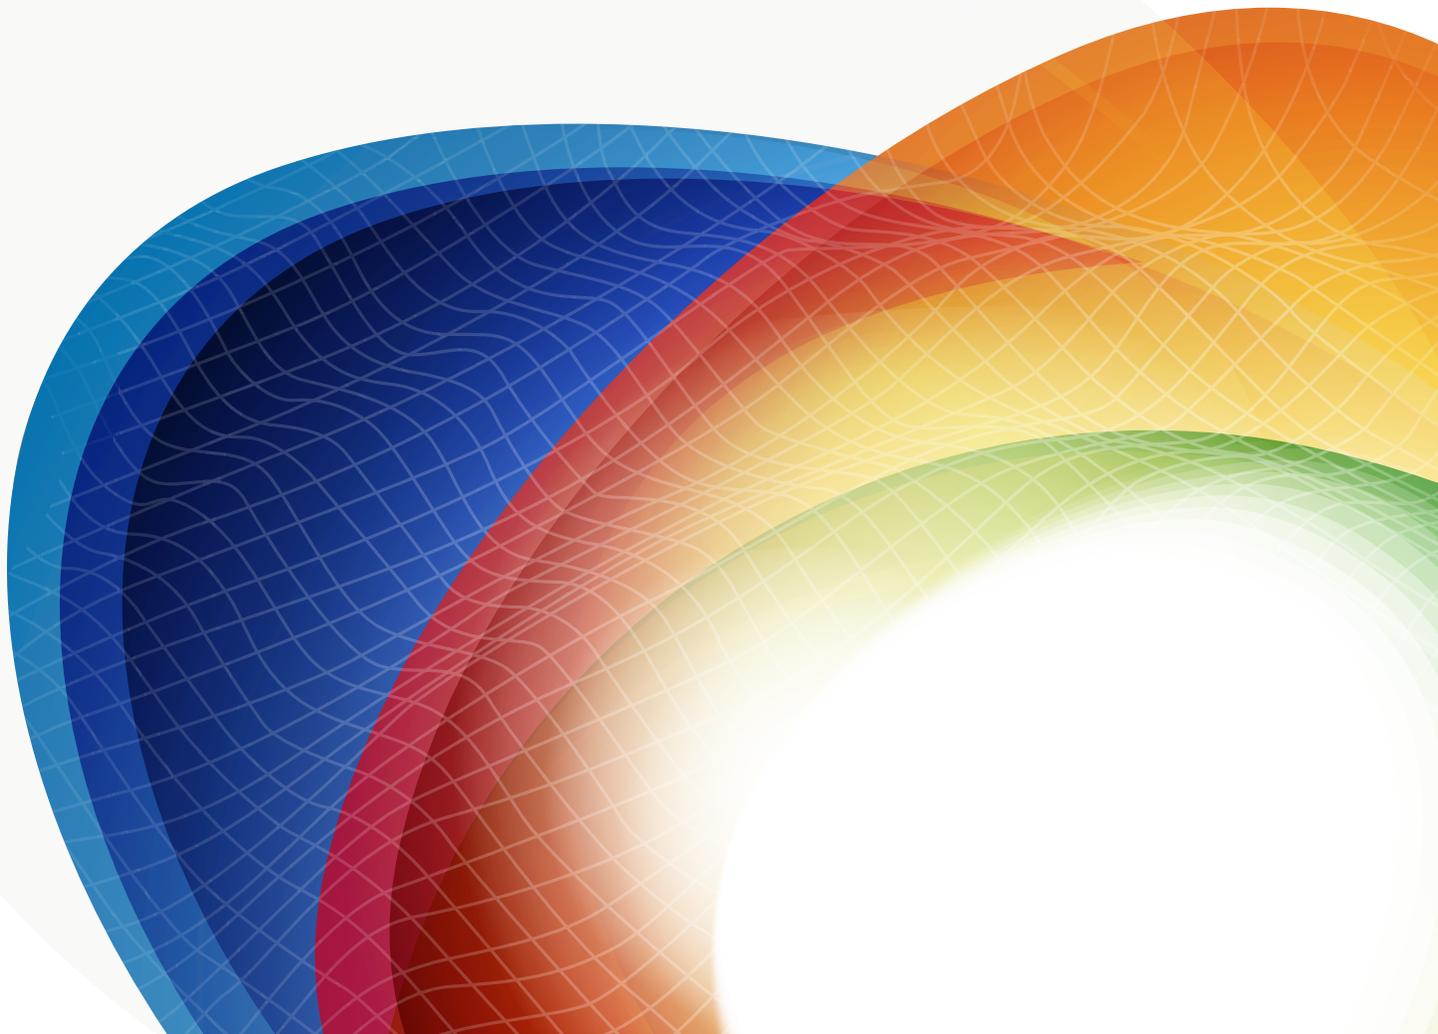




# Software Defined Application Services

Successful management of increased pressure on network and application infrastructure requires flexibility and dynamism across data centers, clouds, and managed environments.

by Lori MacVittie





# Contents

<b>Introduction</b>	<b>3</b>
<hr/>	
<b>The Next Evolution of Application Delivery</b>	<b>5</b>
<hr/>	
<b>F5 Software Defined Application Services</b>	<b>5</b>
The SDAS Platform	7
The Application Service Fabric	8
Application Services	10
<hr/>	
<b>Conclusion</b>	<b>12</b>



# Introduction

Information technology is at a crucial turning point. Enterprise IT departments are under constant pressure to meet user and application demands, aware that cloud deployments offer an easier and faster alternative but often pinned down by legacy deployment models. The problem stems from the inability of those legacy models to adapt to meet expectations for rapid provisioning, continuous delivery, and consistent performance across multiple environments.

Efforts to automate and introduce flexibility continue unabated, with increasing adoption of DevOps and excitement over technologies like software defined networking (SDN) and network functions virtualization (NFV). Yet a significant and critical portion of the data center stack—the application layer—remains deployed in a more traditional model. Enterprises rely on application layer services to address key business and operational concerns such as performance and security, while service providers look to these same services to differentiate offerings across an increasingly fickle subscriber base. It is imperative, then, that these application layer services can be rapidly deployed, easily automated, and ultimately orchestrated within multiple environments and data center models.

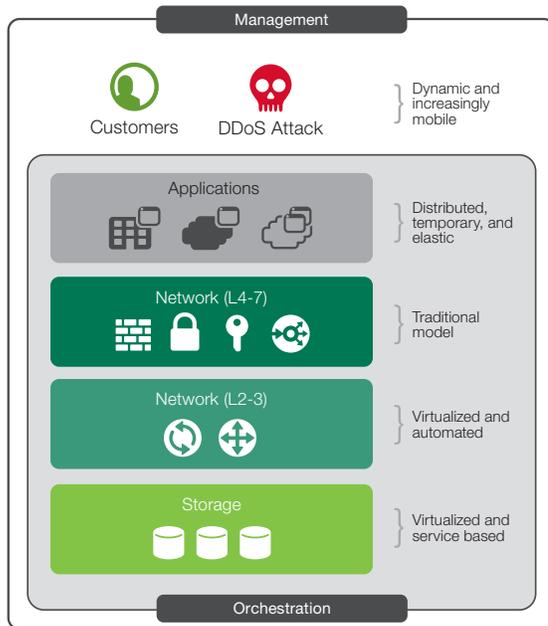


Figure 1: Every layer of the data center stack must become as virtual and dynamic as the applications and users it supports.



When it comes to delivering applications, the challenges of performance, security, and reliability have not changed. What have changed are the environments and conditions under which those challenges must be addressed. In particular, the extension of the data center into cloud environments poses significant obstacles to IT operations trying to maintain consistent policies between data center infrastructure services and those in the cloud. Without architectural parity between the environments, applications may execute without consistent policies for security, performance, and availability. The results include increased risk, unpredictable performance, and loss of control over user satisfaction.

While the changing user landscape has dramatically altered operational policies, applications are also undergoing radical changes in architecture, thanks to virtualization, cloud computing, and the rise of API-based models. These changes directly impact how application services must be provisioned and managed. New services and service provisioning methods have emerged in response, along with new integration models between cloud providers and enterprises, service ecosystems, and service providers. In this rapidly changing environment, the pressure to meet increasing demands often forces IT operations to adopt broadly available, open source solutions despite the complexity those solutions introduce. Such solutions have the advantage of being available now, at little or no cost, which often offsets concern about the longer-term technical debt such options incur. Down the road, however, these organizations can find themselves with a jumble of point solutions that can't communicate with each other, arduous deployment processes, and a manual management nightmare.

It's not just enterprise IT departments feeling the pinch. With pressure to increase revenue per user and a subscriber base that expands by the minute, service providers are seeking new network and application service models that meet their needs for extreme provisioning, value-added services, and competitive differentiation. Traditional service deployment models that use point solutions introduce significant complexity to the Gi network, increasing costs as well as the potential for misconfiguration errors. Furthermore, traditional models simply cannot keep pace with the innovation and growth demanded by the market.

Solving these challenges requires a new model—one that employs greater flexibility in service provisioning and orchestration; one that introduces a new economy of scale for critical application services; and one that enables cross-environment deployments with alacrity and operational efficiency. In this new model, application delivery evolves in alignment with emerging application and network models.



# The Next Evolution of Application Delivery

Application delivery has consistently evolved to meet needs for new delivery models. From a humble beginning comprised mostly of load balancing, application delivery has grown to encompass security, availability, identity verification and access control, cloud services, and performance management. The current excitement around SDN and NFV add programmability and extensibility to this list, since application delivery solutions can rapidly extend network and application services and thus enable rapid service definition and quick response to changing market demands.

How application delivery services are delivered also must change, based on data center automation and the adoption of cloud computing principles. The standard high-availability (HA) pair mitigates failure at the device or instance level, for instance, not the application level. Modern architectures and data center models require a more flexible approach to application services such as availability, one that better aligns with trends toward micro-services and API-based architectures. More broadly, given increased user mobility, an expanding “Internet of Things,” and the reality of HTTP superseding TCP as the de facto transport protocol, service providers and organizations are reevaluating traditional architectural principles to determine how best to move forward with application delivery service provisioning that can keep up with, or at least catch up to, industry trends.

For instance, a key operating principle of virtualization, SDN, and cloud computing is abstraction: the decoupling of services from their underlying infrastructure. Modern architectures favor a foundation of computing and network resources that can be abstracted into a fabric of pure resources that can be sliced, diced, and isolated—and thus more elastically utilized by higher-layer services. A new application service delivery model that also adheres to this principle is necessary to enable the efficiency, responsiveness, and flexibility required in the application service layer of the data center stack.

## F5 Software Defined Application Services

F5® Software Defined Application Services™ (SDAS) is the next-generation model for delivering application services. SDAS takes advantage of F5 innovations in scalability models, programmability, and an intrinsic decoupling of data and control planes to



create a unique application service fabric capable of extending the benefits of F5 application delivery services to all applications, irrespective of location.

SDAS is the first fabric-based application delivery and control system. It enables service injection, consumption, automation, and orchestration across a unified operating framework of pooled resources. SDAS delivers:

- **A fabric-based solution.** F5 ScaleN™ technology powers an elastic, all-active application service fabric that dramatically lowers the cost of delivering application services by increasing utilization and service densities.
- **Automation and orchestration.** Intelligent service automation and orchestration APIs reduce OpEx and fill a critical gap in software defined data center and network architectures. As a result, organizations with SDAS can streamline application deployment and support continuous delivery.
- **A unified operating framework.** A rich, extensible catalog of application services empowers application owners to address performance, security, and availability concerns in cloud, data center, service provider, and managed environments.

The SDAS fabric provides a foundation for building elastic application services. ScaleN technology enables on-demand application and operational scalability at the platform layer. This means the fabric can be deployed on a combination of hardware, software, and virtual form factors—as well as beyond the data center boundary into cloud computing environments—because the platform provides the elasticity and operational consistency needed to scale and manage services in a dynamic environment.

SDAS is comprised of three key components:

- **The application service platform.** The service platform supports programmability of both control and data paths. It is extensible and enables new service creation.
- **The application services fabric.** The fabric provides core services such as scalability, service isolation, multi-tenancy, and integration with the network.
- **Application services.** The heart of SDAS is a rich catalog of services across the application delivery spectrum.

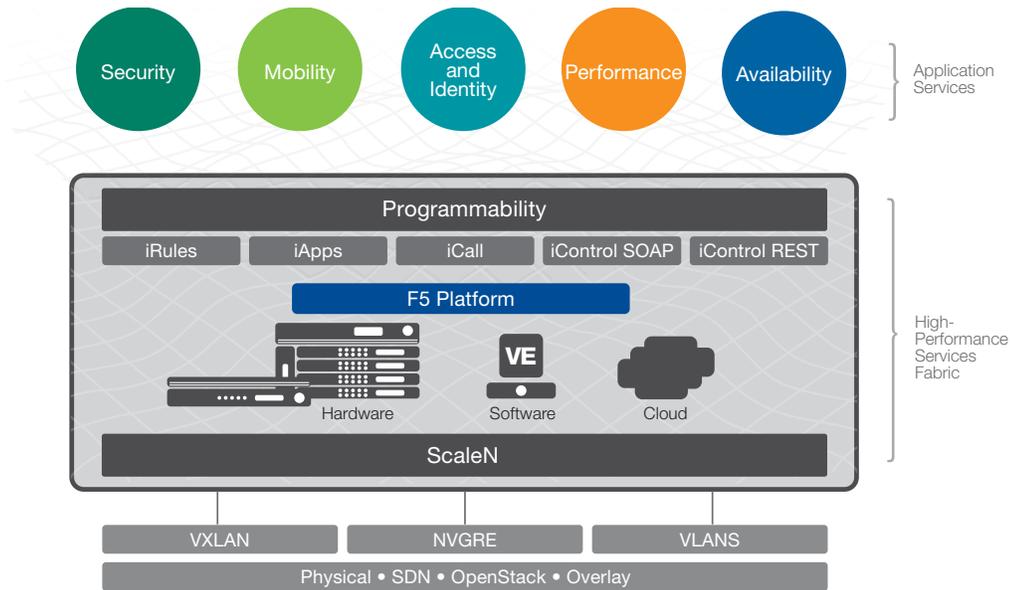


Figure 2: The SDAS architecture comprises an extensible and programmable service platform, a high-performance fabric, and a rich catalog of application services.

## The SDAS Platform

The application service platform enables extensibility and programmability within SDAS. The platform can be deployed in a variety of hardware or software form factors as well as in cloud environments. This provides organizations with flexibility in both scaling and deployment strategies while better balancing capital and operational costs.

The key to the platform is the programmability of the data and control paths. Control path APIs enable integration, automation, and orchestration of tasks such as provisioning, auto-scaling, and migration. Data path programmability provides for rapid response to business and operational conditions as well as new service creation.

### Programmability

Programmability is a key characteristic of SDN models that enables operational extensibility. SDAS embraces programmability in both the data and control paths, ensuring extensibility of the application services as well as the application service fabric.

- **Extensibility.** Functionally, the F5 platform provides a plug-in architecture through which additional services can be added. Over the years F5 has used this architecture to extend application services from load balancing to secure remote



access, application acceleration and optimization, and application delivery firewalling. This platform architecture enables the addition of new application services that address challenges arising from evolving architectures and new protocols, from SPDY and WebSockets to SDN, VXLAN, and NVGRE.

- **Continuous delivery.** Application services are a critical component of the release cycle, and continuous delivery reduces release cycle times and risks through automation of consistent, predictable, and repeatable deployments. The F5 platform includes a number of entry points through which architects and DevOps can programmatically extend and dynamically manage application services. These entry points include the F5 iCall™ technology, F5 iApps® Templates, the F5 iControl® interface, and the F5 iControl REST API. These interfaces, through which F5 SDAS integrates with external orchestration systems, cloud management platforms, and SDN controllers, are supported by an active online community.
- **Rapid system and service provisioning.** Administrators can easily provision and manage fabric resources and application services alike. Organizations and service providers requiring multi-tenancy and self-service capabilities are supported via programmability components as well as through traditional scripts and graphical user and command-line interfaces. Connectors extend provisioning capabilities into the cloud so services can reside in and scale seamlessly across multiple environments.
- **Ecosystem enablement.** SDAS includes a full set of APIs for integration and extensibility of services to support continuous delivery and holistic management of application deployment. Pre-packaged integration with orchestration and cloud management systems such as OpenStack, VMware vCloud, and Amazon EC2 enable rapid implementation for traditional and hybrid cloud deployments.

SDAS optimizes the service orchestration environment, reducing the time to provision services and aligning application services with business and operational requirements. By applying consistent and aligned application services across environments, the F5 service orchestration system can maximize revenue opportunities, enabling organizations to expand into the cloud without compromising on service level agreements or introducing the technical and operational debt incurred by disparate management systems.

## The Application Service Fabric

The adoption of cloud computing has changed expectations with respect to scalability. Elasticity and auto-scaling capabilities are now the norm, particularly for applications with highly volatile or seasonal usage. The application service fabric is an all-active, elastic foundation for SDAS that



addresses the need to move from device-centric to application-focused scalability and exception handling for availability and recovery services.

Based on the F5 platform, the application service fabric scales elastically within the data center via ScaleN as well as into the cloud. ScaleN Device Service Clusters synchronize and share policies across the service fabric, enabling rapid resource provisioning, on the premises or off, and providing a turnkey deployment option for provisioning, licensing, and configuring services regardless of an application's location.

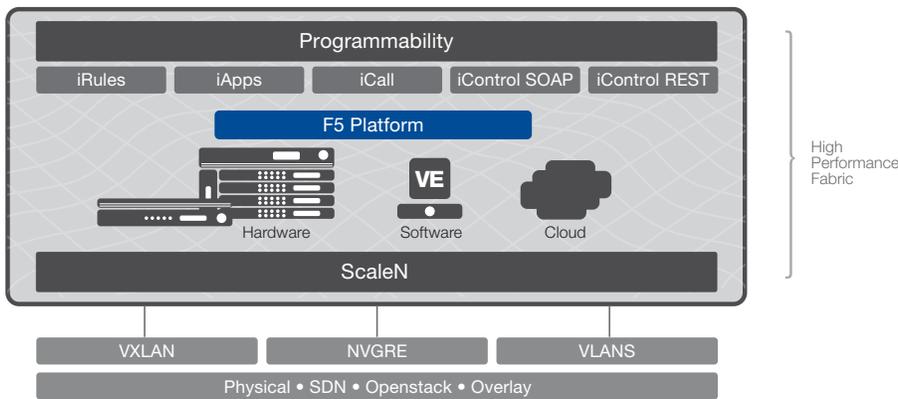


Figure 3: The SDAS fabric interoperates with existing and emerging network models.

The Software Defined Application Services fabric is:

- **Elastic.** All F5 platforms are able to act in concert, irrespective of form factor and across environments. This coordination, which uses ScaleN technology, establishes a trusted, elastically scaled application service fabric that unifies service resources within the data center and into cloud computing environments.
- **All-active.** The application service fabric breaks away from the traditional infrastructure scalability model by eliminating the need for dedicated standby elements, a primary source of operational overhead. This significantly more efficient scalability model better meets the challenges and demands of modern data center architectures.
- **Application-aware.** Emerging data center models emphasize applications, and application delivery models should mirror this focus, enabling context sharing across services and failover at the application layer rather than the device layer. ScaleN provides application failover for the SDAS fabric, ensuring isolation of failure and a greater ability to scale elastically per application or tenant. The application service fabric shares context across services, enabling intelligent service orchestration.



The application service fabric is based on the same economic and operational principles as cloud computing and SDN: a common, shared platform enables economy of scale, rapid extensibility, and operational efficiency by decoupling the foundation from the services. When the underlying infrastructure can be treated as an elastic, operationally consistent fabric, application services can be rapidly and economically provisioned such that all applications, regardless of location, can take advantage of those services.

## Application Services

Application services are deployed atop the SDAS fabric. Comprising a rich array of services spanning security, performance, mobility, and availability, each service scales elastically across the fabric and supports both multi-tenant and role-based management.

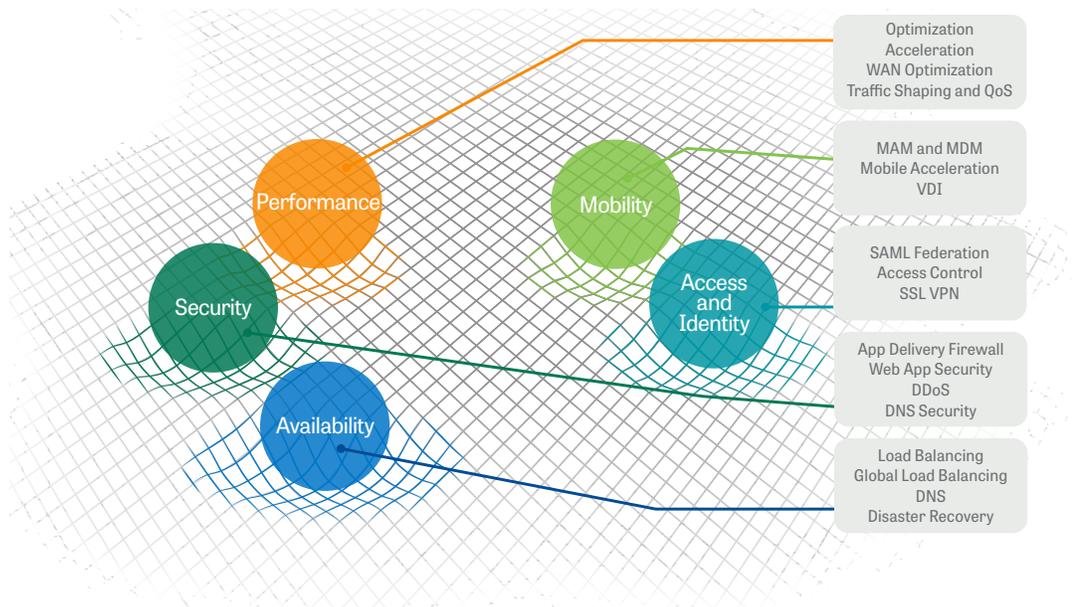


Figure 4: A broad portfolio of application and user services can be rapidly deployed, efficiently managed, and elastically scaled atop the F5 SDAS fabric.

Because they are based on the F5 platform, all SDAS services are inherently:

- **Context-aware.** Each service can deduce a breadth and depth of information about the user, the application, and the network in real time. This enables programmable application and user-aware policies that dramatically improve security, performance, and reliability for users and business stakeholders.
- **Extensible.** Application services can be tailored to meet specific business and operational needs in the data path and based on executable data path logic. A variety of network,



application, and user variables can be used to dynamically apply policies and custom logic, enabling flexible, adaptable policy enforcement as well as custom services.

- **Multi-tenant.** The F5 platform is fully multi-tenant aware, enabling both administrative and network isolation on a per tenant basis. This isolation ensures the security of each tenant by preventing configuration or execution errors from affecting other tenants. Tenants are further insulated through the fault isolation capabilities of the application service fabric.

Services can be orchestrated to ensure the most efficient execution. For example, security services are generally executed before load balancing services to avoid wasting resources on requests that will be rejected due to concerns over origin or content.

F5 application services share a common control plane—the F5 platform. Whether command-line interface (CLI) or GUI, API or templates, all SDAS services reduce operational overhead by ensuring consistent methods of management and integration with F5 and external orchestration and automation systems.

## Conclusion

Applications running across networks and environments encounter a wide range of performance, security, and availability challenges. Addressing these challenges has become more and more complex, particularly when cloud computing enters the picture, often requiring multiple solutions from a variety of vendors with no consistent means of management. This inconsistency is evident in the data center stack, where application delivery—despite its key role in application deployments and user satisfaction—remains fairly static. Cloud computing and SDN have emerged to address specific economic and operational challenges, but they don't efficiently address challenges specific to applications and changing application architectures.

Today, service providers and enterprises need to efficiently provision application services based on the demands and needs of individual subscribers and services. An application that may not need to scale today may still benefit from performance-related services, and subscribers may pay a premium for an optimized mobile experience but not enhanced security. Operations must be able to meet these needs and more with equal alacrity and minimal cost.

F5 Software Defined Application Services delivers economy of scale and extensible, flexible application services with its elastic, fabric-based approach. From the elastic power of the F5 platform, application services can be provisioned rapidly using standardized mechanisms across the data center and into cloud computing environments, reducing the time and costs associated with deploying new applications and architectures.

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

F5 Networks, Inc.  
Corporate Headquarters  
[info@f5.com](mailto:info@f5.com)

F5 Networks  
Asia-Pacific  
[apacinfo@f5.com](mailto:apacinfo@f5.com)

F5 Networks Ltd.  
Europe/Middle-East/Africa  
[emeainfo@f5.com](mailto:emeainfo@f5.com)

F5 Networks  
Japan K.K.  
[f5j-info@f5.com](mailto:f5j-info@f5.com)



**Solutions for an application world.**