



F5 White Paper

Unleashing the True Potential of On-Demand IT

With F5 solutions, users can easily mitigate the complexity and costs of the move toward cloud computing, enabling them to unleash the true potential of on-demand IT.

by KJ (Ken) Salchow, Jr.
Manager, Technical Marketing



Contents

Introduction	3
<hr/>	
Enterprise Virtualization to Cloud Maturity Model	4
<hr/>	
Dynamic Services Model	6
The F5 Dynamic Control Plane	6
<hr/>	
An Ecosystem Approach to Dynamic Services	11
<hr/>	
Conclusion	11



Introduction

If you imagine the flow of anything, be it people, cars, water, electricity or commerce, there are specific points within those flows that provide intelligent processing and control. Typically, these control points exist at strategic points within the flow; discrete boundaries or junctures where multiple flows coalesce for re-distribution. Their purpose is to provide critical flow-control, redirection, intelligent management, and security functionality. The control mechanism might be simple, like a stop sign between two roads, or it might be complex like power-distribution and control systems in modern high-rise buildings. It might be physical, virtual, or procedural. Regardless of form, without such a control point the efficient, safe, and uninterrupted flow of any system would be impossible.

At the same time, the origination and destination of traffic rarely stays static over time. As flows change, the strategic points of control often cannot adapt dynamically and thus become impediments to efficient flow. As roadways—where a simple stop sign once promoted controlled movement of traffic—become highways, the previously adequate control point becomes a dangerous bottleneck or worse. Often, the process of changing the type or function of a control point to meet the new demands can be long, complicated, and costly.

Today, IT faces a variety of changes that require control points that can adapt dynamically. Server virtualization, cloud computing, and the staggering explosion of mobile users and devices make the traditional static control points obsolete. Mobile users continually change devices, networks, and access methods. At the same time, server virtualization and cloud computing alternatives constantly alter the destinations of that traffic.

Fortunately for IT professionals, it is now possible to design a network using strategic points of control that dynamically adapt to changes in flow, and thus constantly provide each user with optimal access to their applications and data.



Enterprise Virtualization to Cloud Maturity Model

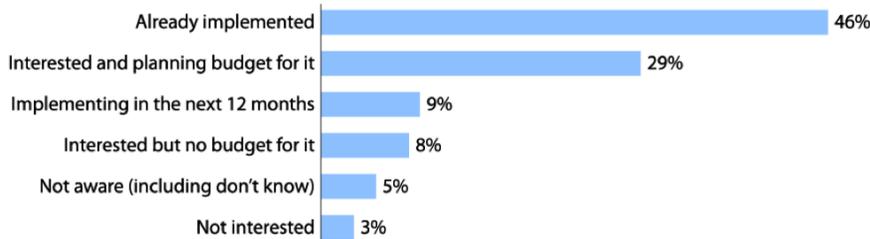
Enterprise IT is in transition. In a Q3 2008 survey of 118 decision-makers at North American and European enterprises, a total of only 8% stated that they were “Not Interested” (3%) or were “Not Aware (including Didn’t Know)” (5%) concerning x86 virtualization, when asked about their highest level of awareness or interest on the topic¹. At the same time, few, if any, have moved all their business critical applications to the cloud environment. Most organization’s solutions exist between two extremes of static, dedicated servers or on-demand IT. Some organizations have solutions in multiple points within the cloud maturity model, depending on the particular application or business unit. Some might never evolve to incorporate external cloud provider services.



April 2009 “Capacity Management In A Virtual World”

In Enterprises, X86 Virtualization Is Mainstream

“What is your company’s highest level of awareness of or interest in x86 server virtualization?”

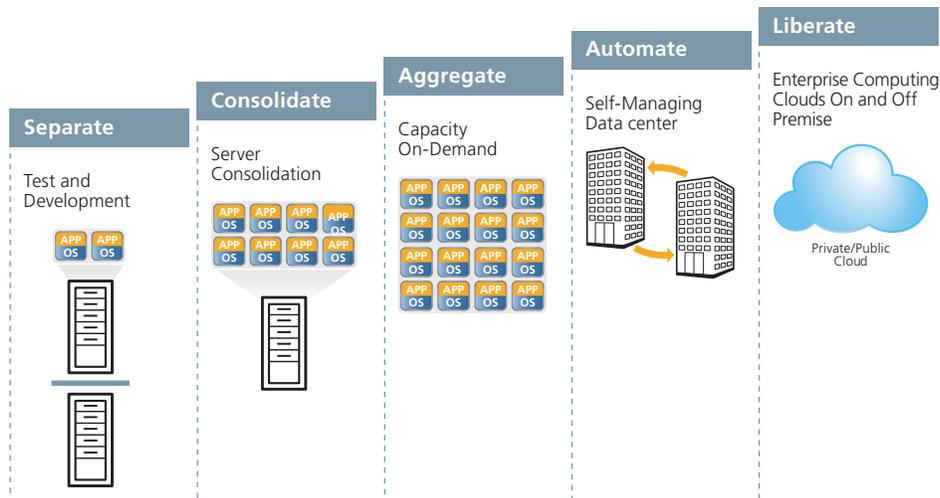


Base: 118 x86 server virtualization decision-makers at North American and European enterprises

Source: Enterprise And SMB Hardware Survey, North America And Europe, Q3 2008

© 2009, Forrester Research, Inc. Reproduction Prohibited

In enterprises, x86 virtualization is mainstream.



Enterprise virtualization to cloud maturity model

Regardless of why an organization has started to implement server virtualization—green IT, data center consolidation or other cost reduction—most have discovered that actual benefits have not lived up to the hype. While initial goals might be met, unanticipated side effects often reduce the realized benefit. One good example of this is server management. If an organization does not have a well-designed and efficient server management practice, the process of virtualizing physical servers makes server management exponentially more difficult. The same is true of patching, security, and a host of other areas.

In addition, the mobility of users combined with the new mobility of applications, has added a layer of complexity to ensuring that users can reach their applications and data in a consistent and usable manner. As organizations continue from enterprise virtualization to cloud maturity, one outcome remains constant; the challenges of the past are amplified and magnified. The increased level of complexity, coupled with reduced control over applications and data that leave the corporate data center, quickly brings many organizations to a point of diminished returns long before they achieve the goal of on-demand IT.

Businesses are looking for a way to maintain the flexibility of these technologies while containing the inherent complexity. Businesses want freedom from the bounds of their own data centers, but need to maintain the same levels of control in security, optimization, and management that exist with static, dedicated servers. They want the ability to understand the myriad aspects of application and data flows in relation to users. To do so, businesses need new types of control points in the data center that dynamically and intelligently adapt to change.



Dynamic Services Model

Organizations need a dynamic services model providing reusable services that understand context and can provide control of information flows regardless of application, virtualization, user, device, platform, or location. This model does not necessarily call for replacing existing investments and is not likely provided by any single vendor. The dynamic services model is an evolutionary way to instrument existing systems—application delivery, security, optimization, virtualization, and management—creating a holistic ecosystem and eliciting more value from each existing component as it integrates with that ecosystem.

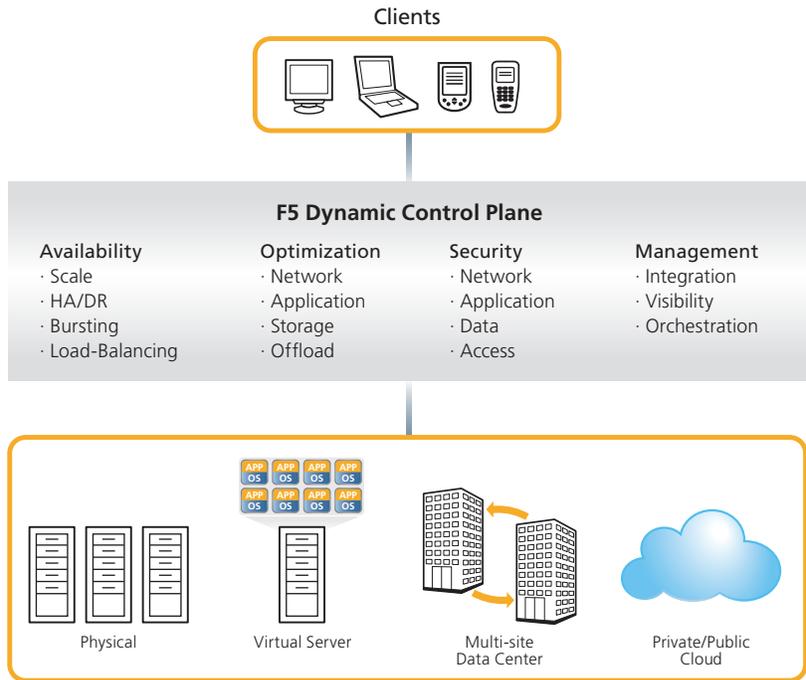
One of the key characteristics of this model is to deploy strategic points of control to create a coordinated mediation layer between the users, the applications, and the data they access. In this manner, you can characterize the dynamic services model as the final abstraction between the users and the applications. This layer presents a consistent interface and set of services that can be used for all applications and data regardless of their current location. This layer uses an intelligent understanding of context—“who” is accessing “what” from “where” and “why”—to determine the optimal connection between users, applications, and data. Finally, it provides a deep understanding of context to inform the underlying application and data elements about the current resource requirements and to instruct the infrastructure to adapt in real-time to ensure optimized application and data access.

The dynamic services model describes the ultimate strategic point of control—the combination of a stop sign, metered stoplight, express lane, and hi-speed bypass. Capable of dynamically changing its nature and the nature of the surrounding infrastructure based on the unique needs of each individual car, this strategic point of control can instantly turn a two-lane road into a superhighway.

The F5 Dynamic Control Plane

F5 refers to its implementation of dynamic services as the dynamic control plane architectureⁱⁱ.

The dynamic control plane architecture is the extension of F5’s entire suite of application and data delivery solutions, integrating and coordinating with the virtualization, management, security, and application components often already present within the infrastructure.



F5's dynamic control plane architecture

The best way to understand the power of the dynamic control plane is through real world examples.

Cloud-Based Load Balancing and Traffic Management

Many organizations have deployed application delivery solutions for their business-critical applications. These deployments have enabled high availability, offloaded security and optimization, and have promoted simplified management and control of these applications. However, when an organization wants to move those applications—either temporarily to handle outages or spikes in demand or permanently—to external hosting providers and/or cloud providers, how does it proceed if the service provider does not offer the same application delivery solution? Organizations must either find a way to translate the existing application delivery configurations to the provider's solutions—and hope that they provide the same functionality and performance—or they must purchase and pre-position their own Application Delivery Controllers in the provider network. This significantly reduces the flexibility of the solution and increases the cost of moving from one provider to another. It is easier to re-route traffic if you do not have to build the detour from scratch every time.

F5 provides the flexibility to deploy application delivery solutions on a choice of platforms, including a virtual platform that has the same functionality and configuration as its hardware-based counterparts. This gives F5 customers the



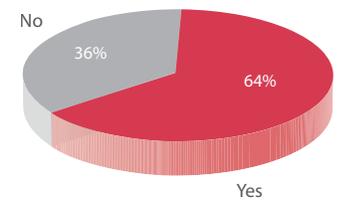
flexibility and ease of control that is required with a dynamic services model. Customers can deploy F5 Application Delivery Controllers in dedicated data centers and, at a moment's notice, deploy a virtual version with the exact same functionality for security, optimization, and availability in any service provider's network. The organization's unique policies and configurations can be migrated with the application as it moves to the service provider. With parity of functionality, organizations can guarantee the exact same user experience, while significantly reducing costs associated with changing cloud providers. And because all of the application delivery functions are integrated, new virtual devices are easily and automatically provisioned with the F5 BIG-IP® Global Traffic Manager™ (GTM).

Long Distance Migration of Live Applications

A similar problem exists with the applications themselves. While hypervisor vendors like VMware provide tools to enable organizations to move applications from one piece of physical hardware to another without affecting users, these tools have limited capability to move applications between data centers—or to external service providers—even with dedicated connections. Ideally, organizations want the flexibility of moving applications at any time, to any location without affecting the application's users. However, WAN or Internet connections are often too slow to accomplish this feat. Even if the tools themselves do not fail because of excessive delay, the time it takes to move the application makes it nearly impossible to continue servicing user requests. This is particularly true if the application has a back-end disk image that must also be migrated; while the application might move to the new location, the delay caused by having to access its data store in a remote data center would bring the user experience to a stand-still. Lastly, moving the application does not mean that the user's requests automatically transfer to the new location. Traffic redirection is required to avoid lost connections and poor performance.

F5's dynamic control plane integrates multiple F5 technologies, as well as those from other vendors. The integration between F5 and solutions from VMware, HP, and Microsoft enables real-time dynamic provisioning of services. If the orchestration engine recognizes that compute resources are below set thresholds or using excessive resources, it can dynamically move the application to hardware with additional resources and notify the F5 solution of the new or added application instances. A live application migration between data centers or to a cloud service provider might occur because an application instance lacks sufficient resources to run locally.

Inter-Data Center VM Movement Critical for Cloud Providers



Source: TechValidate
TVID: A6D-C70-242



When long distance live migration occurs, the dynamic control plane applies the appropriate services to ensure consistent delivery. F5 devices dynamically create optimized and secure tunnels between devices. These tunnels serve multiple functions. The tunnel can significantly reduce the time required to move a machine and its disk image between data centersⁱⁱⁱ, making it possible to use the same migration tools whether the target is in the next rack or around the world. The tunnel can also reroute existing user connections to the new application location with minimal impact on user performance.

This process is possible only because the F5 dynamic control plane applies services intelligently, and is inherently open to collaboration and integration with the greater ecosystem.

Universal Access Control and Acceleration

Access control and acceleration of applications present significant challenges when applications are released from their physical server bounds. Managing authentication, authorization, and auditing (AAA) services while ensuring enforcement of enterprise security policies on individual machines is difficult enough, but virtualizing the applications and potentially distributing them to multiple data centers and service provider networks can leave an organization unsure of their security posture. Remote networks, especially those not entirely under the control of the organization, are particularly difficult to manage. These networks might not have the same enforcement capabilities and might require the replication of sensitive user credential databases outside the corporate data center. One solution is simply to route all user requests to the corporate data center for authentication and then redirect users to remote locations for application access; however, this can make the corporate data center a bottleneck and affect user accessibility. Finally, without centralized control, an organization may still lose the ability to properly audit and log user access, which can cause significant compliance issues if the organization cannot accurately identify who accessed sensitive information, when it was accessed, and for what purpose (application). If there are many roads leading to the same destination, traffic laws must be enforced equally on all paths or traffic will quickly learn how to circumvent the law; yet, enforcement of the law should only slow law-breakers, not all traffic.

F5 provides the unique ability to combine various services to solve these types of problems. Both the BIG-IP® Edge Gateway™ and the BIG-IP® Access Policy Manager™ (APM) provide integrated solutions to enable organizations to handle AAA services in a controlled and simplified manner. These components allow



organizations to integrate client integrity verification, per-application authentication to various back-end authentication solutions, and single-sign-on functionality. In addition to being able to create optimized and secure tunnels back to the corporate data center, organizations can now deploy consistent security enforcement in multiple locations and still leave sensitive authentication databases in the corporate data center. BIG-IP GTM, using geolocation data, ensures that users connect to the best access point. That access point can then provide AAA services—using resources in the corporate data center.

If the optimal user access point and requested application reside in different locations, the dynamic control plane leverages secure, optimized tunnels to connect users to applications without sacrificing the implementation of corporate security or the user experience.

The dynamic control plane provides simplified, consistent control to the organization, applying the right services at the right time to optimize connections between users, applications, and data. All information flows obey the same set of equally enforced rules.

Consistent, Real-Time User Performance Measurement

The increased mobility of users and numerous connections used to access corporate applications and data make it difficult to monitor the true user experience accurately. While organizations like Gomez provide sophisticated solutions to manipulate the data stream for real-time performance measurement, not every application is customizable enough to allow the addition of these measurement components. This can be true even when an organization has full control of its applications and data. When applications escape the confines of the corporate data center this inability to track performance is especially problematic. If you cannot measure the real user experience, it is difficult, if not impossible, to ensure that co-location and cloud peering-point facilities are maintaining the SLAs they promise. Ideally, all applications, regardless of type or location, should be monitored in real-time, with the resulting knowledge used to optimize their delivery dynamically.

As a strategic point of control point located between applications and users, the BIG-IP® Local Traffic Manager™ (LTM) offers a universal tagging point to automatically inject tracking components into all traffic to all applications—not just those that can be customized. And, since BIG-IP LTM is available in physical and virtual instances, it can provide this function beyond the confines of the corporate data center.



An Ecosystem Approach to Dynamic Services

It is important to stress that the dynamic services model is an ecosystem approach; it is not a single vendor solution or a point product. The programmatic and procedural resources provided to integrate, coordinate, and collaborate with other ecosystem elements are defining characteristics of F5's dynamic control plane.

This is evident in the illustrated solutions with VMware and Gomez, but these are only a few examples of the F5 dynamic control plane solving real-world problems today. F5 maintains formal relationships with leading technology providers including Dell, HP, IBM, ExtraHop, Infoblox, NetApp, CA, Symantec, webMethods, Secure Computing, RSA, WhiteHat Security, Splunk, TrendMicro, ByteMobile, Microsoft, and many others. These relationships include tested and documented integration with F5's solutions and, as such, can be thought of as an extension of the dynamic control plane architecture. The architecture is open, allowing partners and customers to create and deploy their own custom applications. Virtually any service requiring consistent implementation across servers, data centers, or the cloud, can be instrumented within the dynamic control plane and extended ecosystem.

More importantly, the dynamic services model is designed to provide a holistic solution as organizations adapt from virtualization to cloud. And, it grows with the organization's efforts, enabling movement from one level to the next. The dynamic provisioning services are certainly important in a single data center, but become critical in a multi-data center environment. Long-distance migration of live applications is important in multi-data center solutions, but becomes critical with the introduction of cloud services. The model supports and enables the IT organization in rapidly responding to change, and consequently the value of the model increases with the amount and intensity of change.

Conclusion

Regardless of where an IT organization falls within the enterprise virtualization to cloud maturity model, its leaders must respond to an endless variety of demands. Rapid response to change requires a holistic solution that enables flexibility and simplicity without losing context and control. The dynamic services model describes the ultimate strategic point of control—a stop sign, metered stoplight, express lane, and hi-speed bypass combined. Capable of dynamically changing its nature and the

White Paper

Unleashing the True Potential of On-Demand IT

nature of the infrastructure surrounding it based on the needs of each individual, this strategic point of control can instantly turn a two-lane road into a superhighway.

The dynamic services model presents a new perspective on IT architecture, providing the foundation for reusable services that understand context and can provide control regardless of application, virtualization, user, device, platform, or location. It adds value to traditional, static, dedicated-server application infrastructure, enables more modern virtual implementations, and becomes a critical requirement of cloud services deployment.

F5's dynamic control plane is a real-world embodiment of the dynamic services model. F5 solutions provide IT organizations the ability to grow beyond static data centers, manage movement and migration of their applications and data down the hall or around the world, support demanding and ever-increasing mobile users, offer centralized and consistent security enforcement and global, dynamic measurement of real-time application performance. With this approach, F5 customers are better equipped to mitigate the complexity and costs of the move toward cloud computing, enabling them to unleash the true potential of on-demand IT.

ⁱ "Capacity Management in A Virtual World," Forrester Research, Inc., April 2009

ⁱⁱ See "[Controlling the Cloud](#)" for a detailed examination of F5's dynamic control plane

ⁱⁱⁱ [VMware Deployment Guide](#)

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

